
CERN

CERNBox

EOS

Cloud Storage for Science

Massimo Lamanna
On behalf of the CERN IT STorage group (*)

CERN – Soleil visit (22-JUN-2016)

CERN

(*) Collaboration with CERN EP-SFT (P. Mato, D. Piparo and E. Tejedor)

Outline

- Storage at CERN
  - Short recap
- EOS
  - Disk solution for LHC
- CERNBox
  - Cloud storage services
    - based on EOS and ownCloud

Storage at CERN

WLCG computing

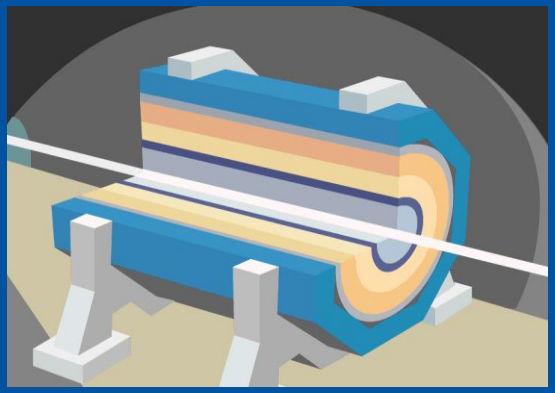

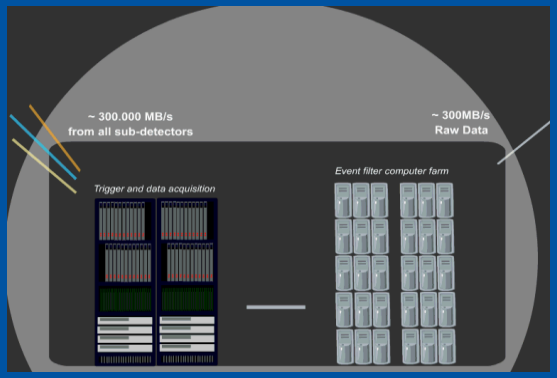

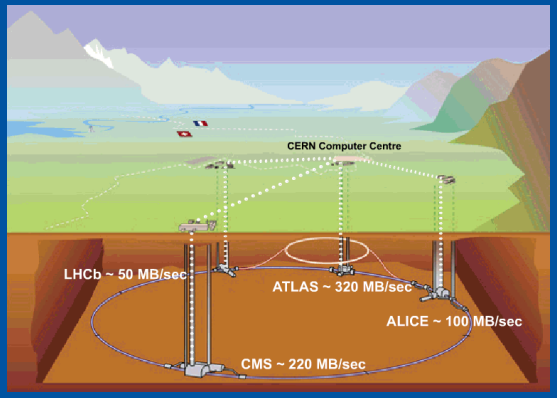

- Reconstruction use case
  - Goals: data quality and immediate access for analysis
  - Organised activity dominated by heavy processing and replication
- Analysis use case
  - Goals: extract physics quantities (discovery)
  - Individual activities dominated by event selection and sharing

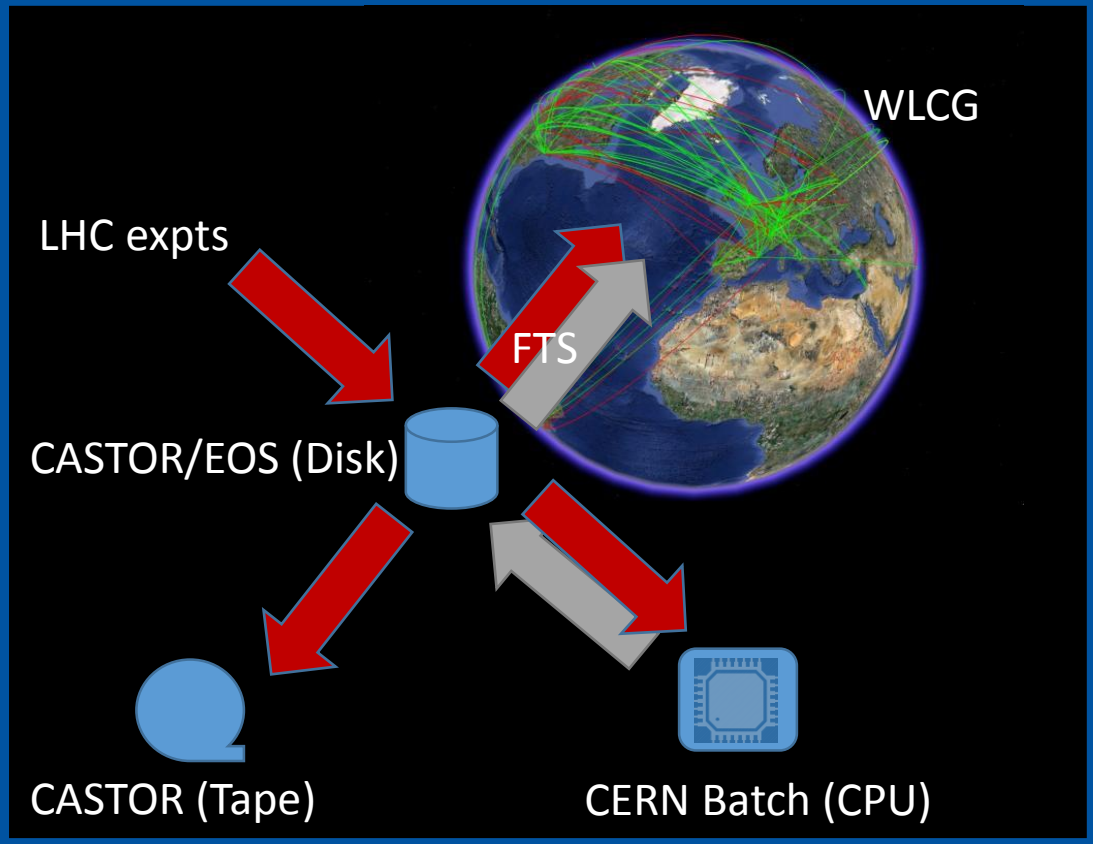

IT ST data services

AFS

- Home directories (user quota up to 100 GB)
- O(1) PB
- Backup

CASTOR

- Tape archive with disk staging area
- O(100) PB archive

EOS

- Disk farm
- O(100) PB disk space

EOS/CERNBox

- User space (sync client; initial user quota 1 TB)
- AFS use cases for new/large projects
- Backup user space

Ceph

- Virtualised infrastructure foundation
- CERN has a seat in the Ceph Advisory Board
- Block storage (OpenStack), Object store, …

FILER

- High-performance NFS servers

TSM, HADOOP, …

EOS

CERNBox

I LIKE HOW WE'VE HAD THE INTERNET FOR DECADES, YET "SENDING FILES" IS SOMETHING EARLY ADOPTERS ARE STILL FIGURING OUT HOW TO DO.

Randall Munroe (http://xkcd.com)

(LHC) Data Taking 2012 → 2015

CASTOR
CERN Advanced STORage manager

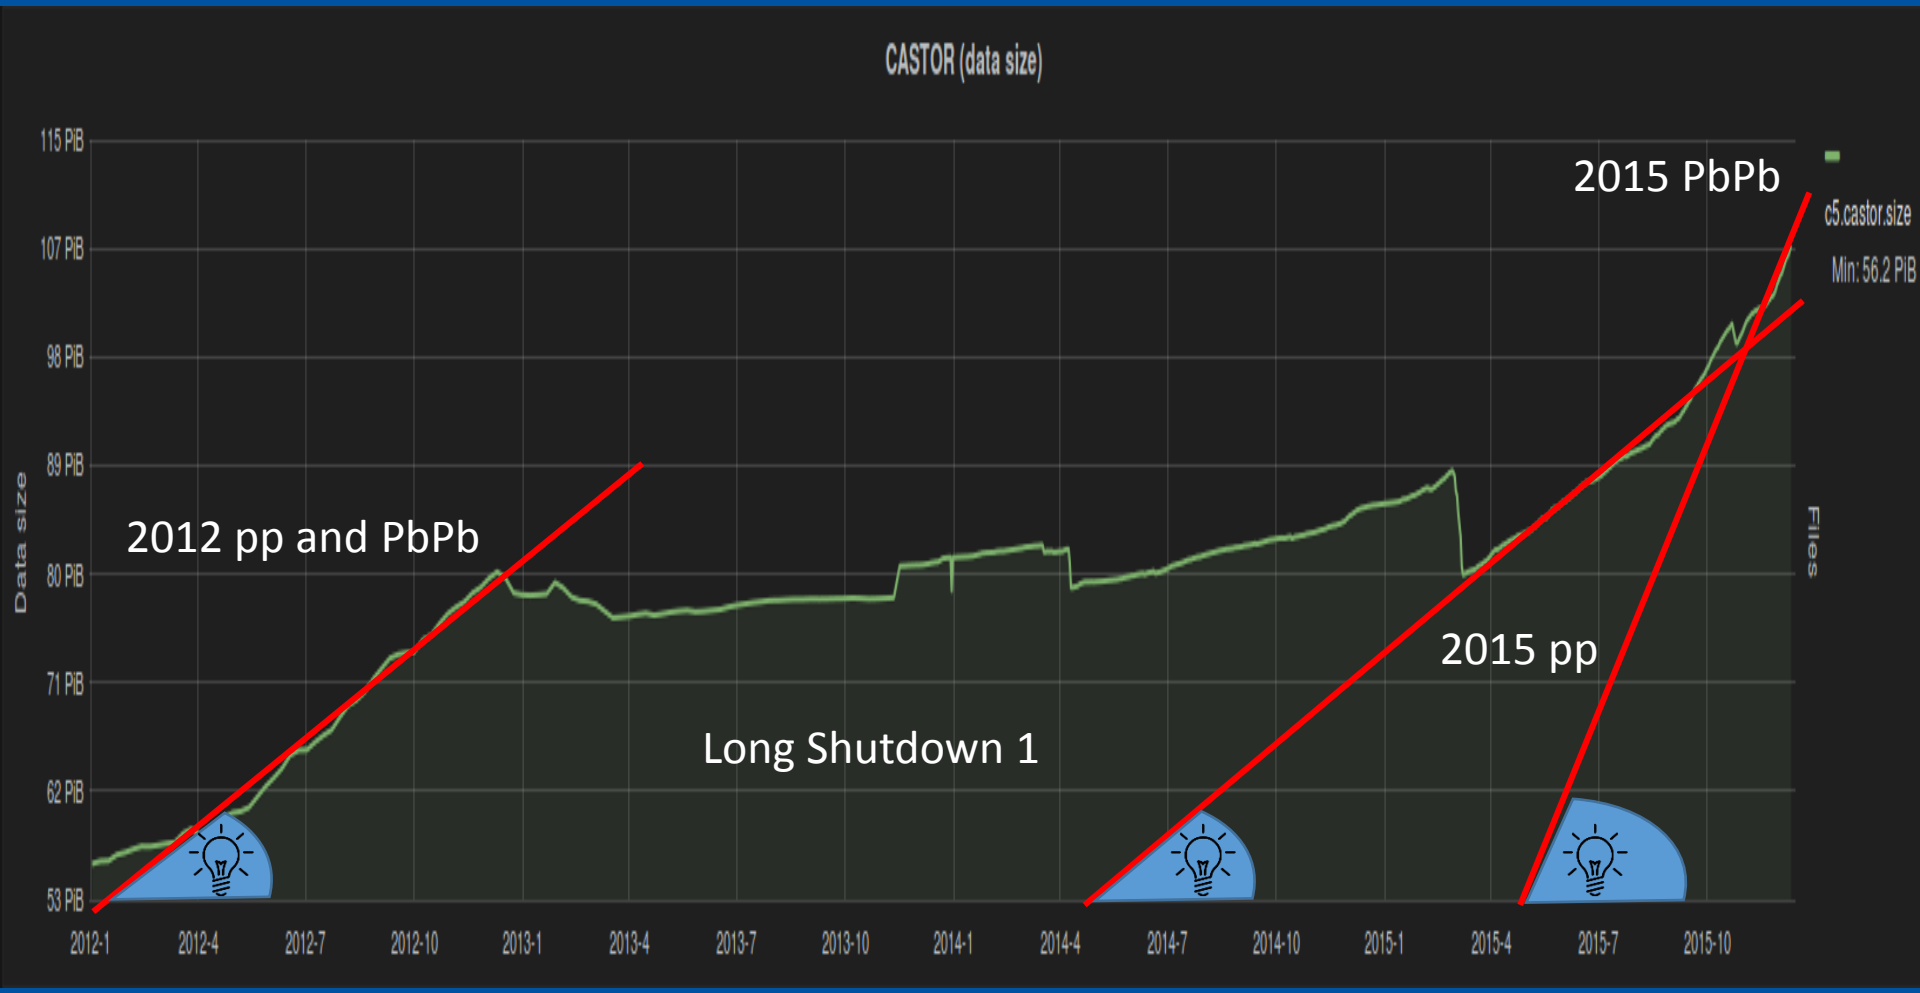

Last week record: 2 PB in a week

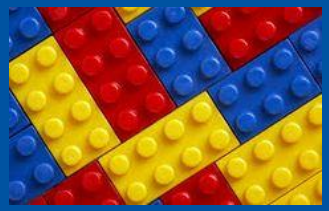

Commodity / Uniformity

- Uniform HW across services
  - Same building blocks for AFS, CASTOR, Ceph and EOS
  - AFS:CASTOR:EOS 1:10:70 ratio
  - Heterogeneity essentially from technology evolution
    - Yearly cycle; new models with best “price per GB”
  - Recently took over NetApp NFS services
    - Now on the same hardware than EOS & co.

- Main building block
  - ¼ of quad system (1 node with dual CPU, 64 GB RAM, 10Gbit/s NIC)
    - OpenStack hypervisor use the same hardware
  - 2 24-disk trays per node
  - 24 6-TB disks per tray
    - 55 PB delivery in 2016Q2
  - → 200 TB building block

EOS

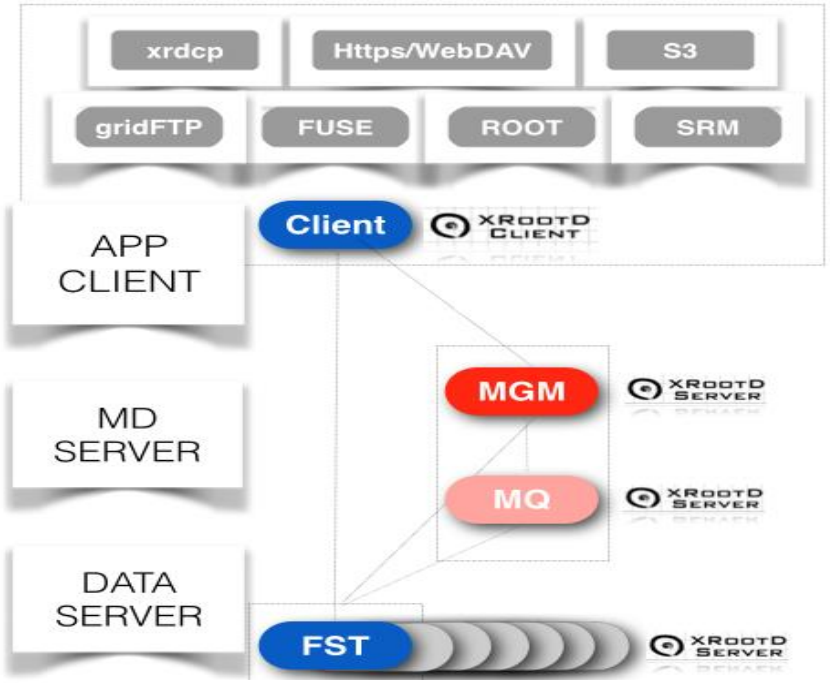

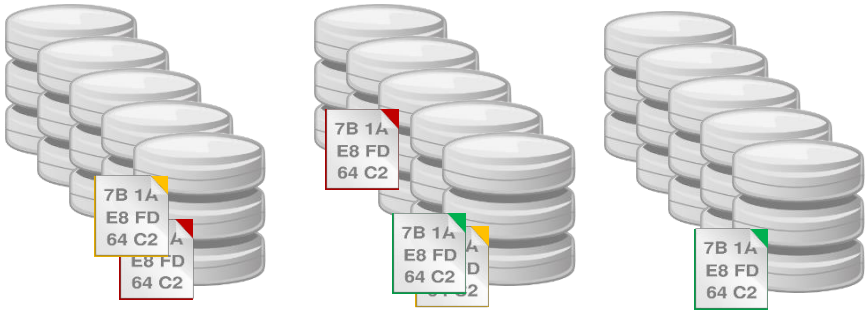

- EOS: Large disk farms for physics and beyond
  - Towards 190 PB JBOD
- Developed in CERN/IT
- Original goal
  - Large scale (PBs for 100s/1000s independent scientists) analysis of LHC data
- Strategic points
  - Distill 20+ years of experience in HEP/data management
  - Ultra-fast name space
  - Optimised protocols
  - Arbitrary level of data durability via cross-node file replication or RAIN using commodity hardware (complete control on the JBOD layer)
- Status
  - Open to non-physics use cases
  - Large number of protocols available
  - "All" client platform
    - Red Hat Linux + OSX
    - (More Linux + Windows)
  - Catalogue scale up:
    - 10^8-10^9 → 10^{10}-10^{11}

Our “20-ms-large” computer centre

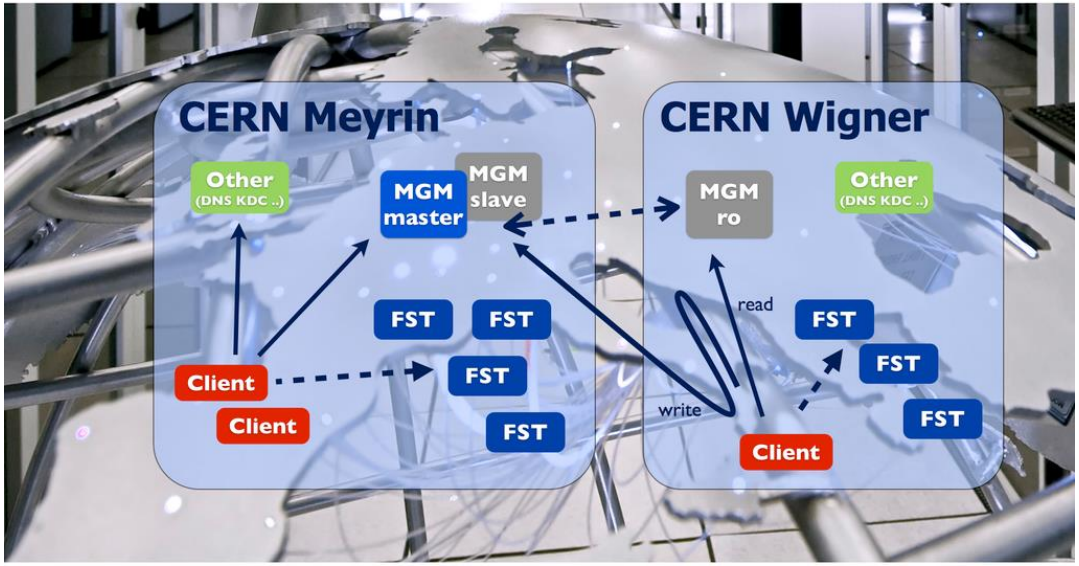

Autonomic,
Locality,
Disaster recovery/
business continuity

MGM = NameSpace/Metadata FST = Disk servers

EOS (April 2016)

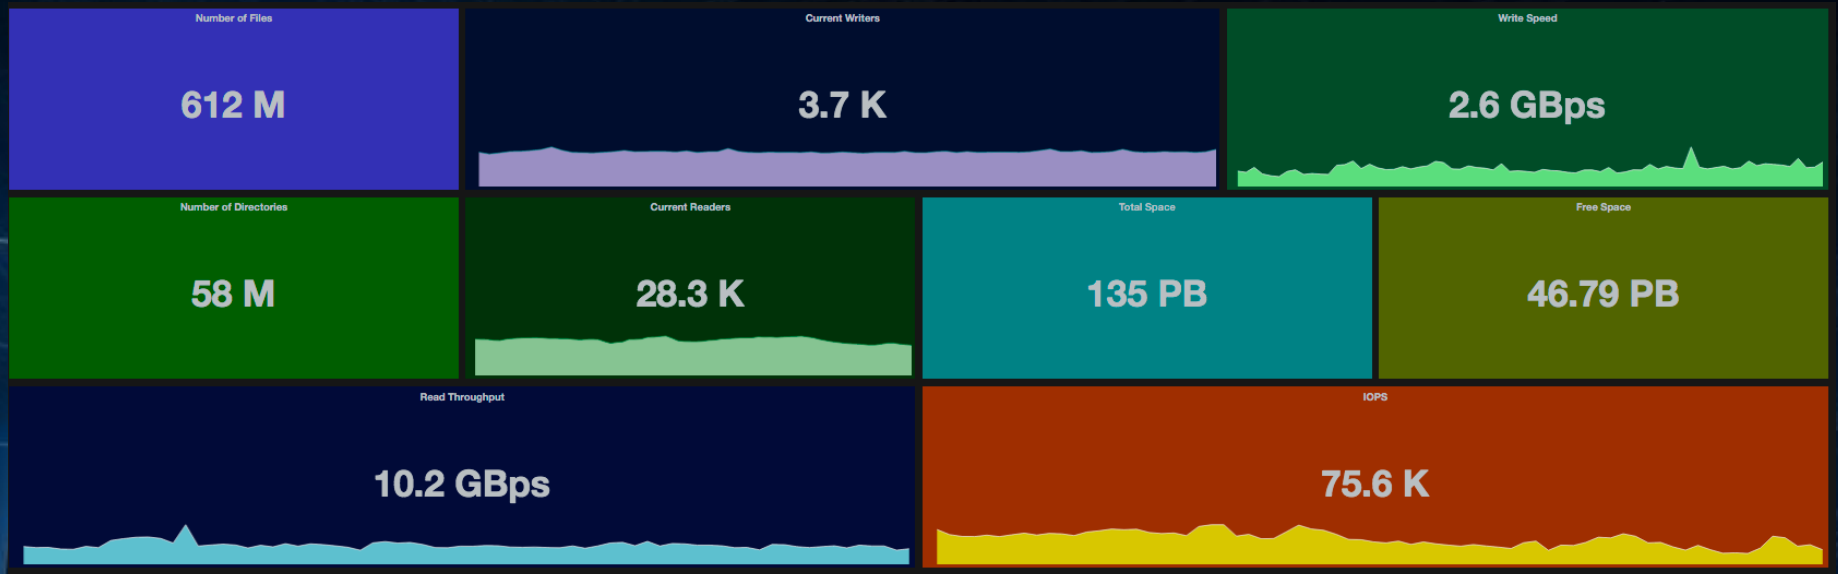

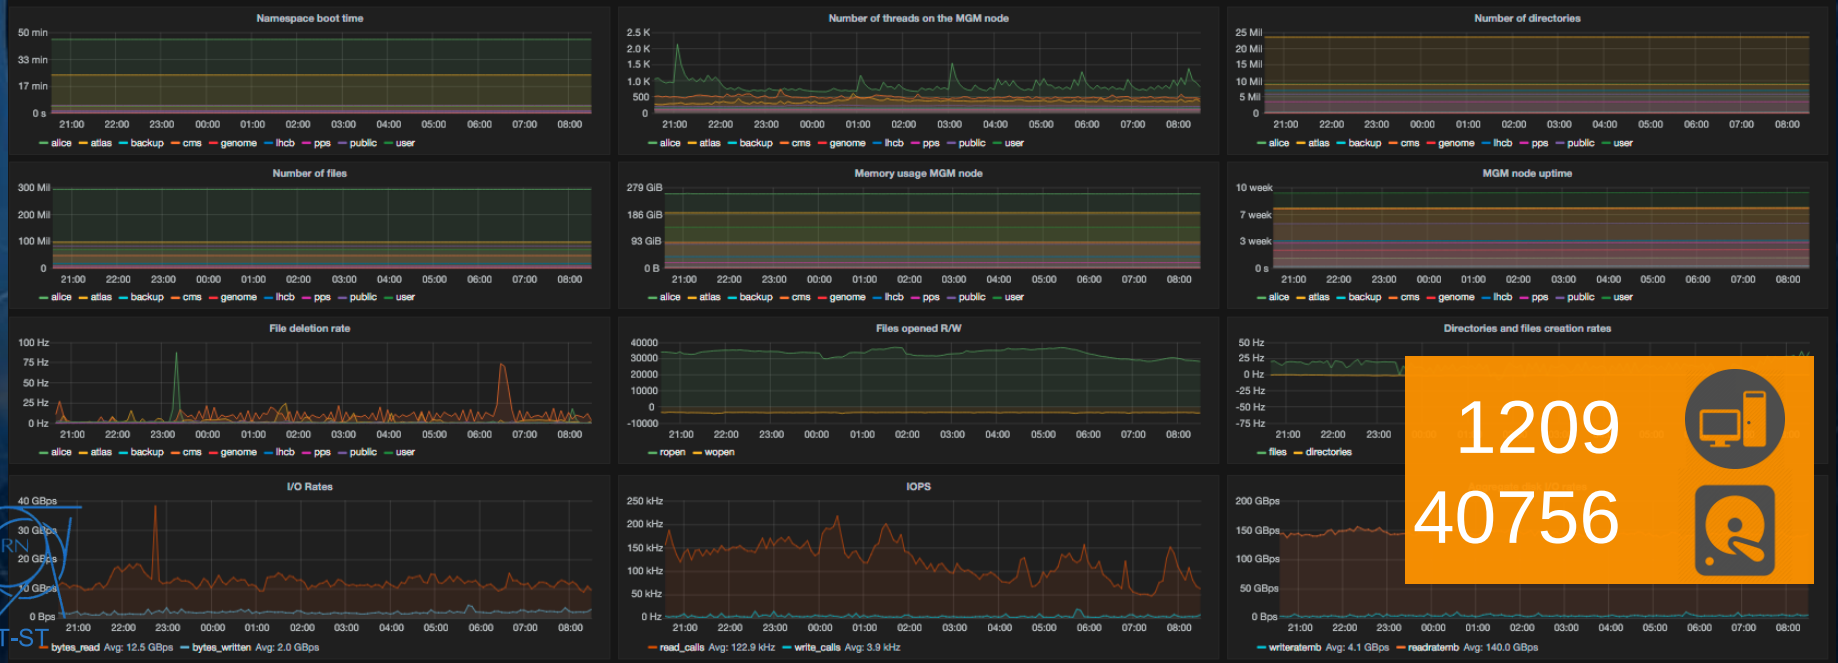

1209

40756

CERN IT-ST

CERNBox

Rationale behind cloud storage

- Existing professional usage of public cloud storage (Dropbox) at CERN
  - CERN (IT)
    - Harmonically extend our service portfolio
  - Today users expect:
    - Cross-platform, easy install and setup (BYOD)
    - Modern web interfaces
    - Mobile access

  - Selected ownCloud as starting point
    - Open software; data on premise
    - General public use and deployments in edu/research

  - **Scale up in size**: interesting challenge (leverage on EOS)
  - **Extend in the use-case phase space**: innovation

CERN

What is CERNBox ?

CERNBox provides a cloud synchronisation service

- Synchronise files (data at CERN) and offline data access
- Easy way to share with other users
- All major platforms supported
- Based on ownCloud integrated with EOS
  - Available for all CERN users (1TB/user initial quota)

Much more than a Dropbox™ replacement!

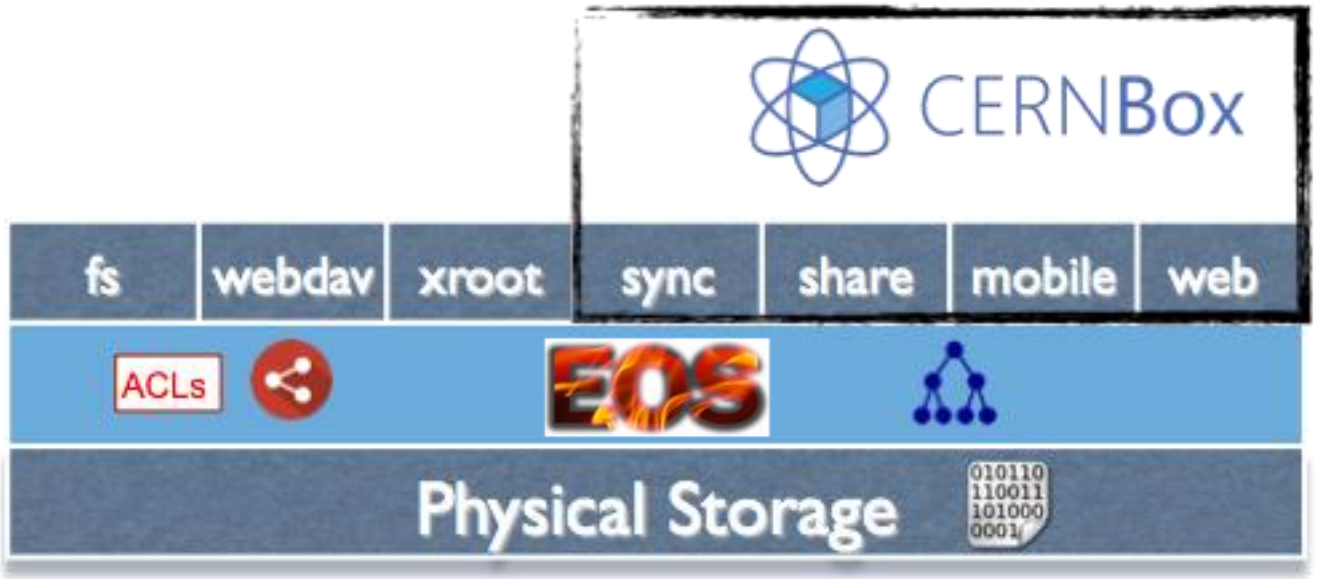

Access Methods: Sync

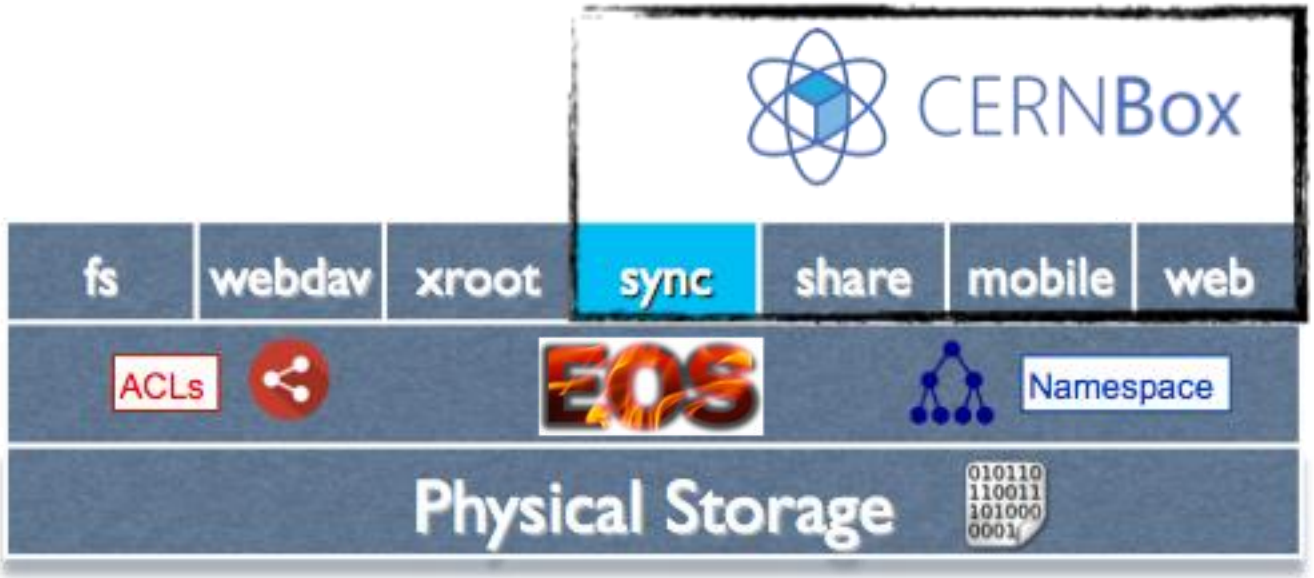

Access Methods: Sharing

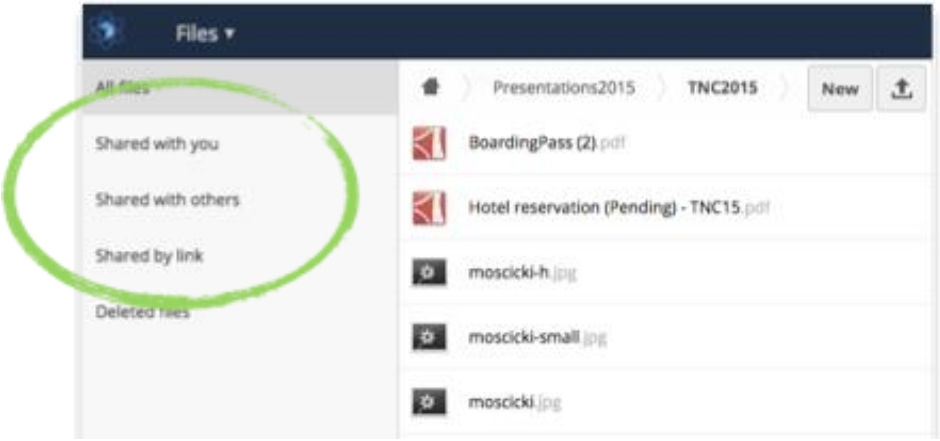

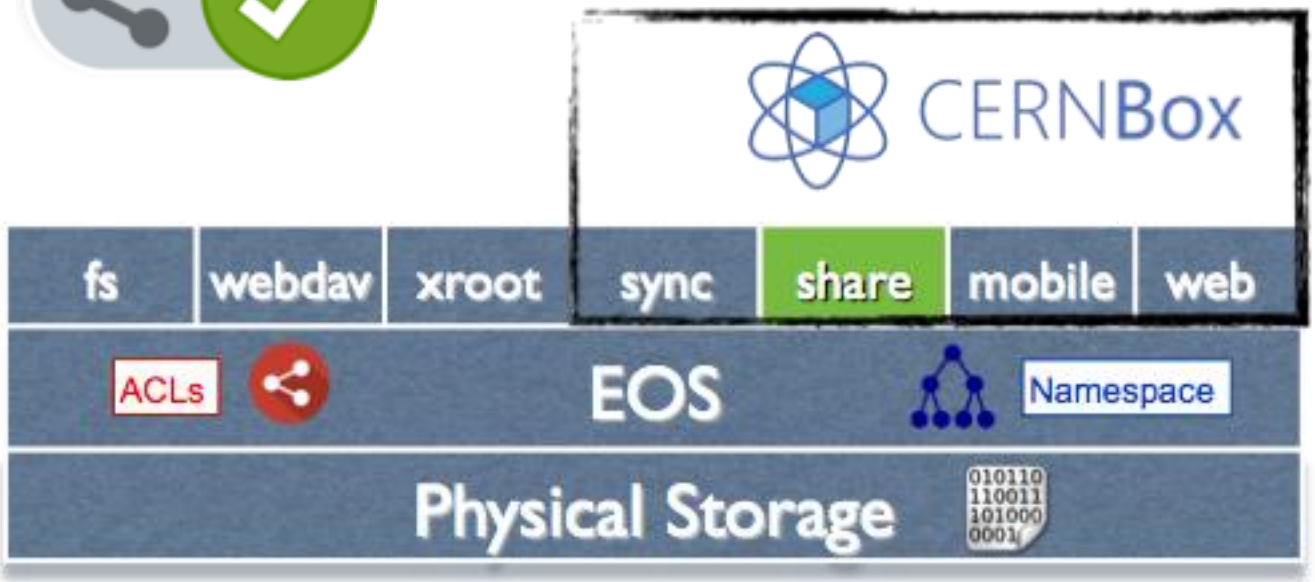

Access Methods: Mobile & Web

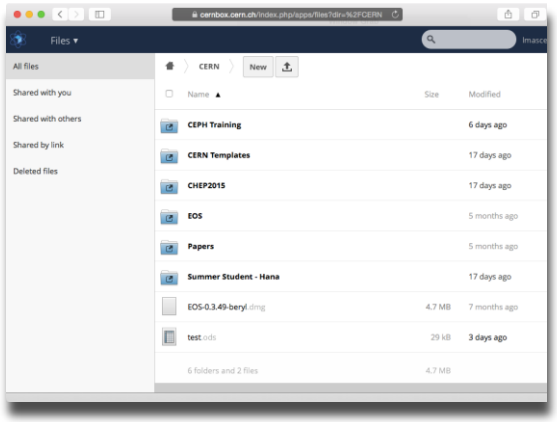

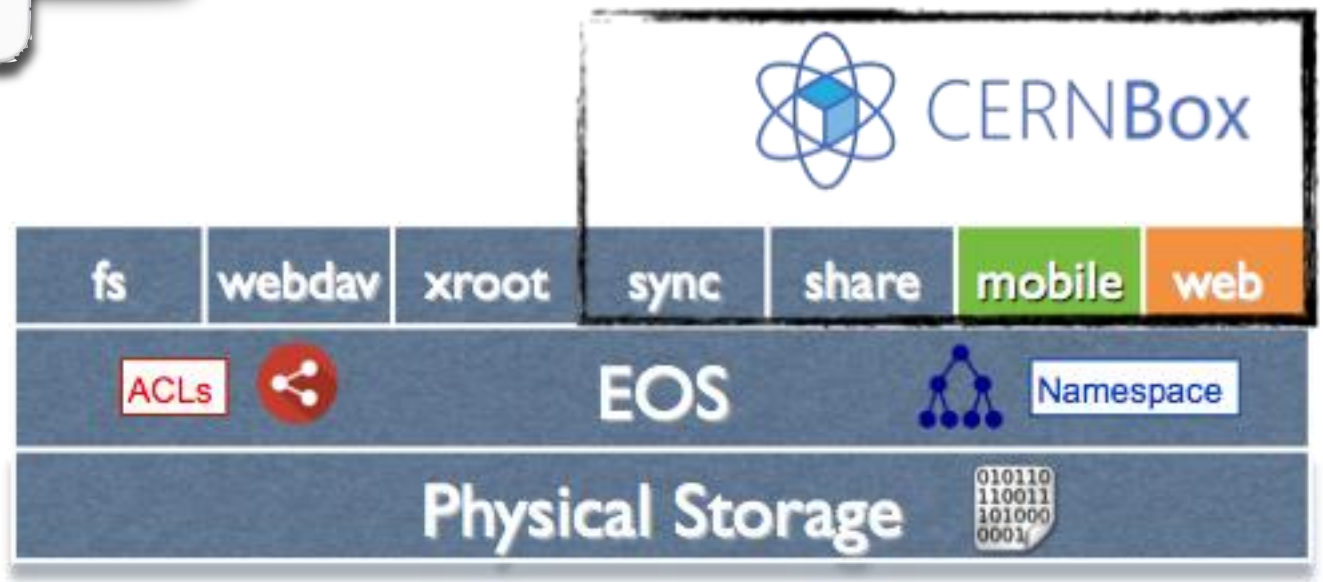

Access Methods: WebDAV

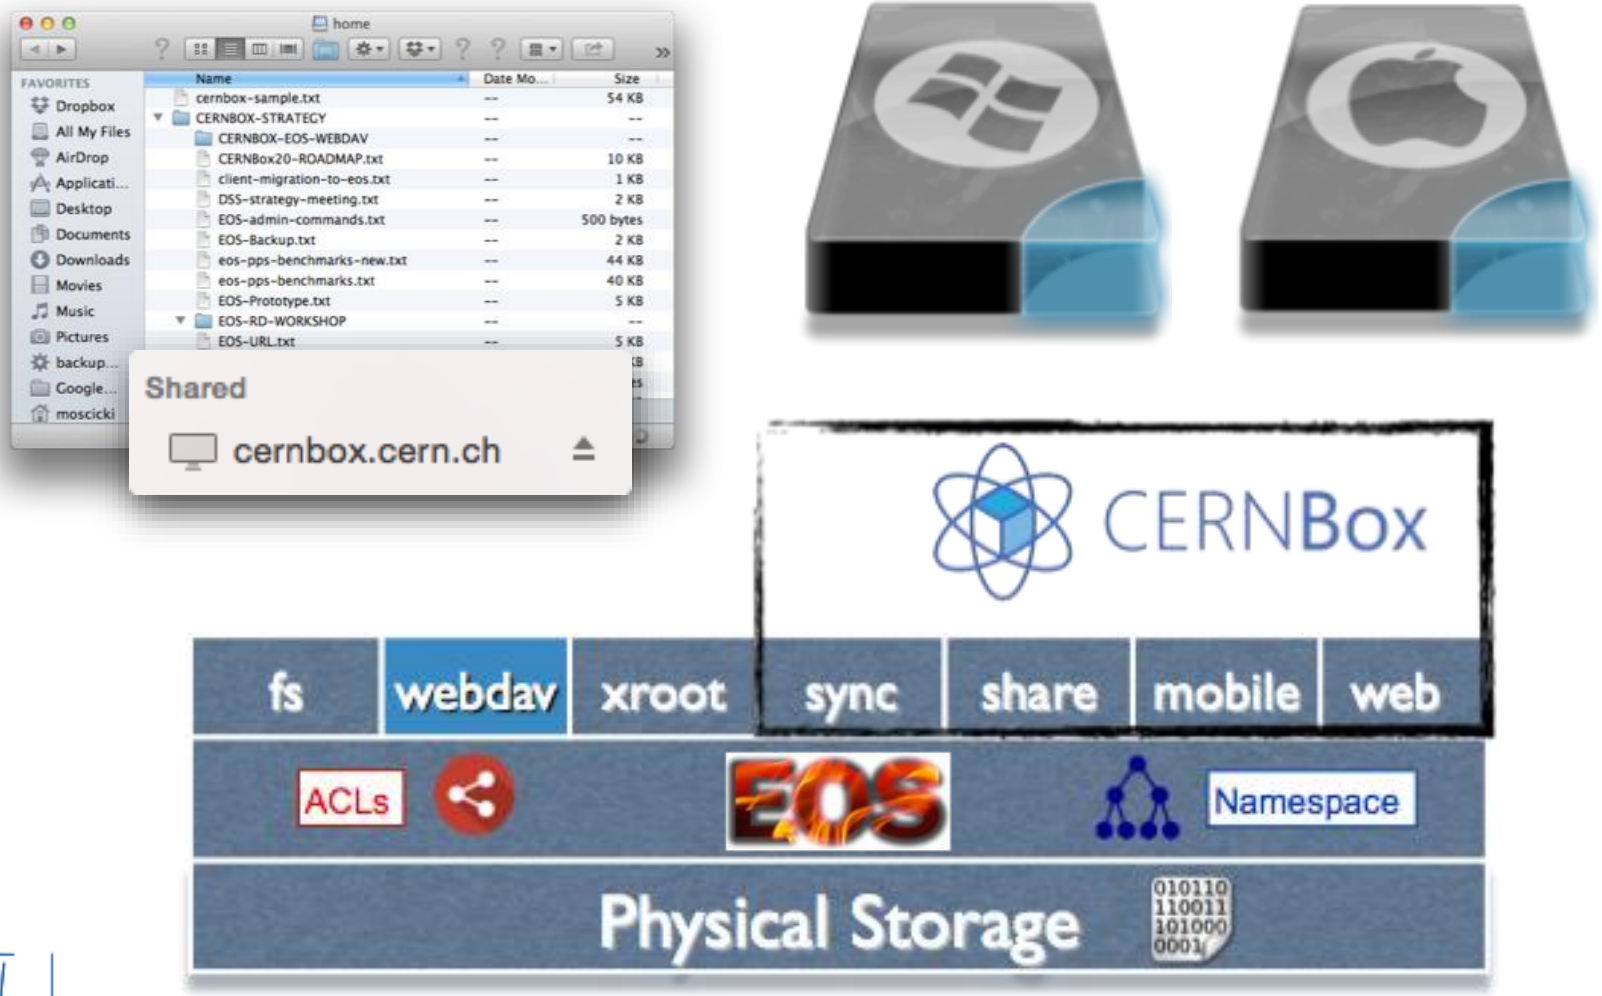

Access Methods: FUSE


```
[lmascett@lxplus2015 ~]#
[lmascett@lxplus2015 ~]# df -H -t fuse
Filesystem      Size  Used Avail Use% Mounted on
eosuser         506T   70T  437T  14% /eos/user
eosatlas         36P   17P   20P  45% /eos/atlas
eosalice         20P   11P  8.5P  57% /eos/alice
eoscms           28P   14P   15P  49% /eos/cms
eoslhcb          13P  7.6P  4.6P  63% /eos/lhcb
eospublic        16P  5.8P   11P  36% /eos/public
[lmascett@lxplus2015 ~]#
[lmascett@lxplus2015 ~]# ls -lc /eos/user/l/lmascett/
total 6644
drwx------. 1 lmascett c3       5 Dec 10 15:58 CERN
drwx------. 1 lmascett c3       0 Jan 26 18:18 debug
drwx------. 1 lmascett c3       0 Dec 11 09:43 download
drwx------. 1 lmascett c3       0 Oct 31 18:24 pdf
drwx------. 1 lmascett c3       1 Dec 11 09:44 personal
drwx------. 1 lmascett c3       8 Dec 10 12:11 pictures
```

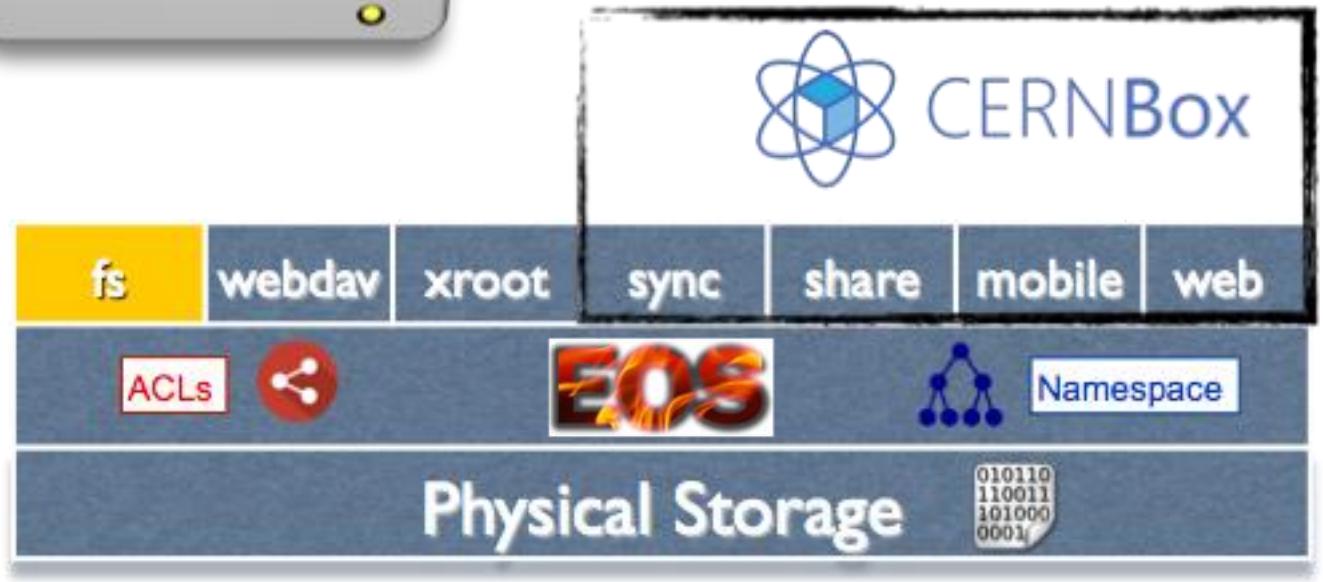

Optimised access

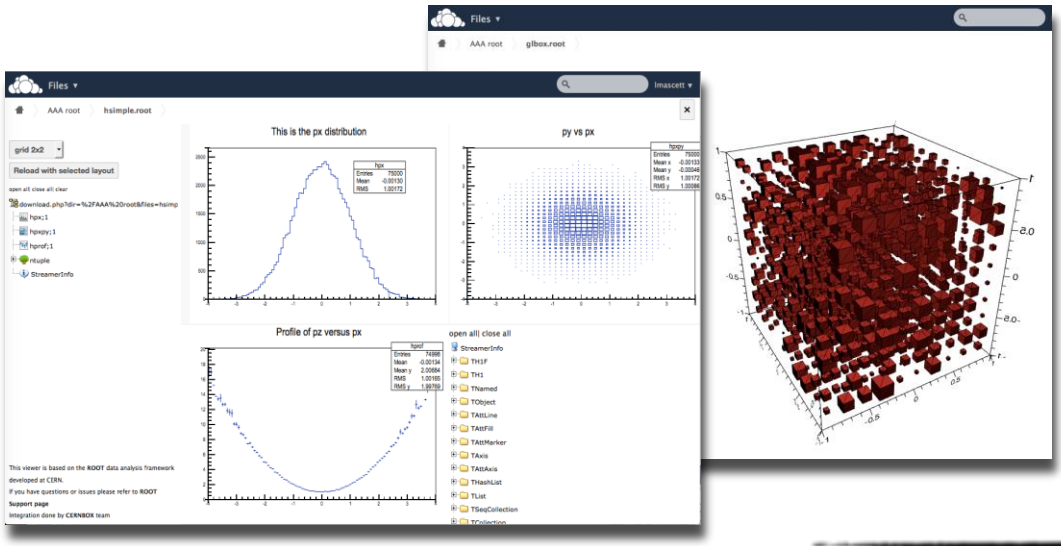

Embedded ROOT viewer
in CERNBox browser

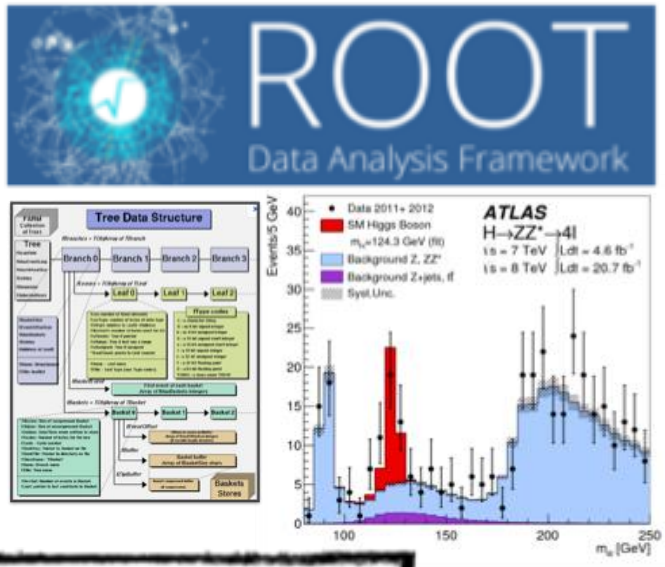

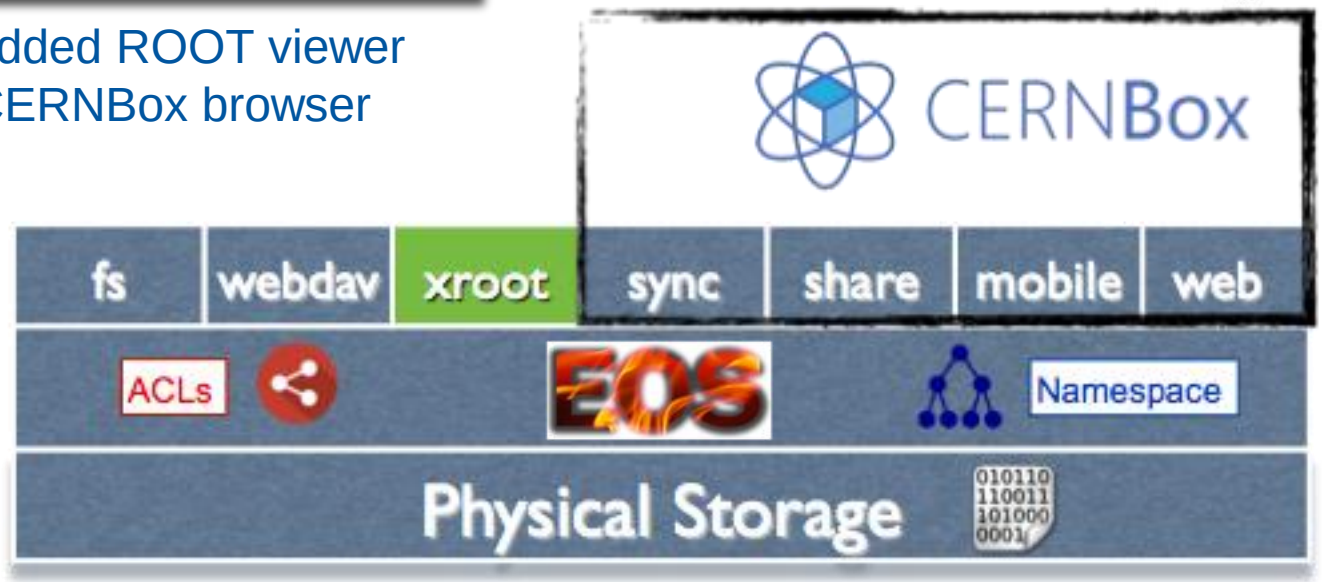

CERNBox service

Users	4337
# files	55 Million
# dirs	7.2 Million
Quota	1TB/user
Used Space	104 TB
Deployed Space	1.3 PB

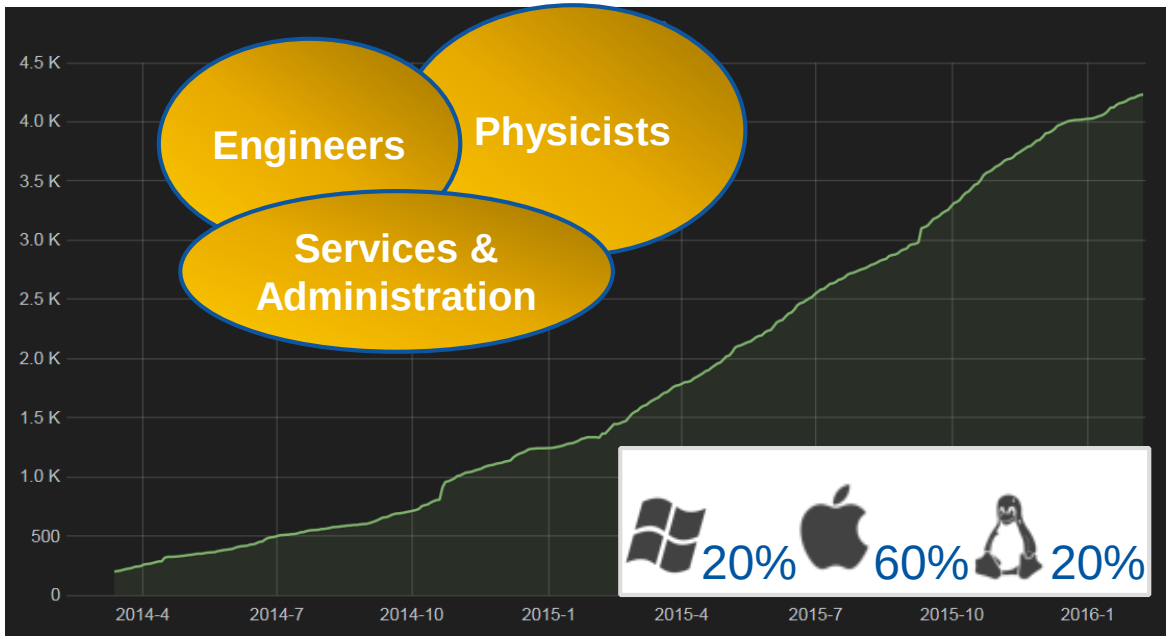

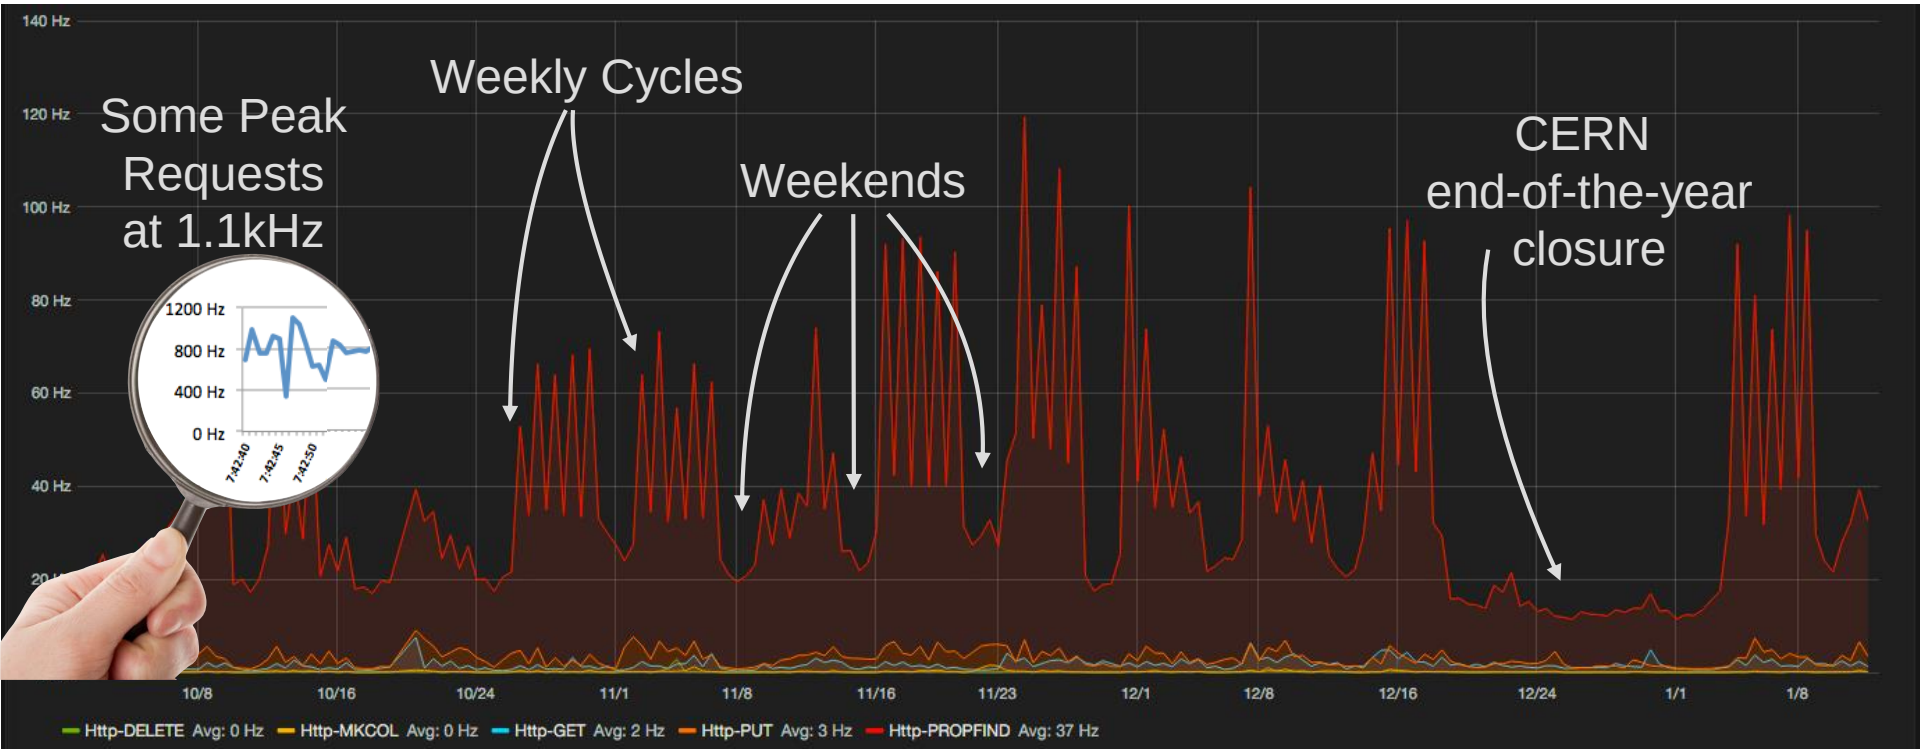

Dec 2015: Geolocation Active Users

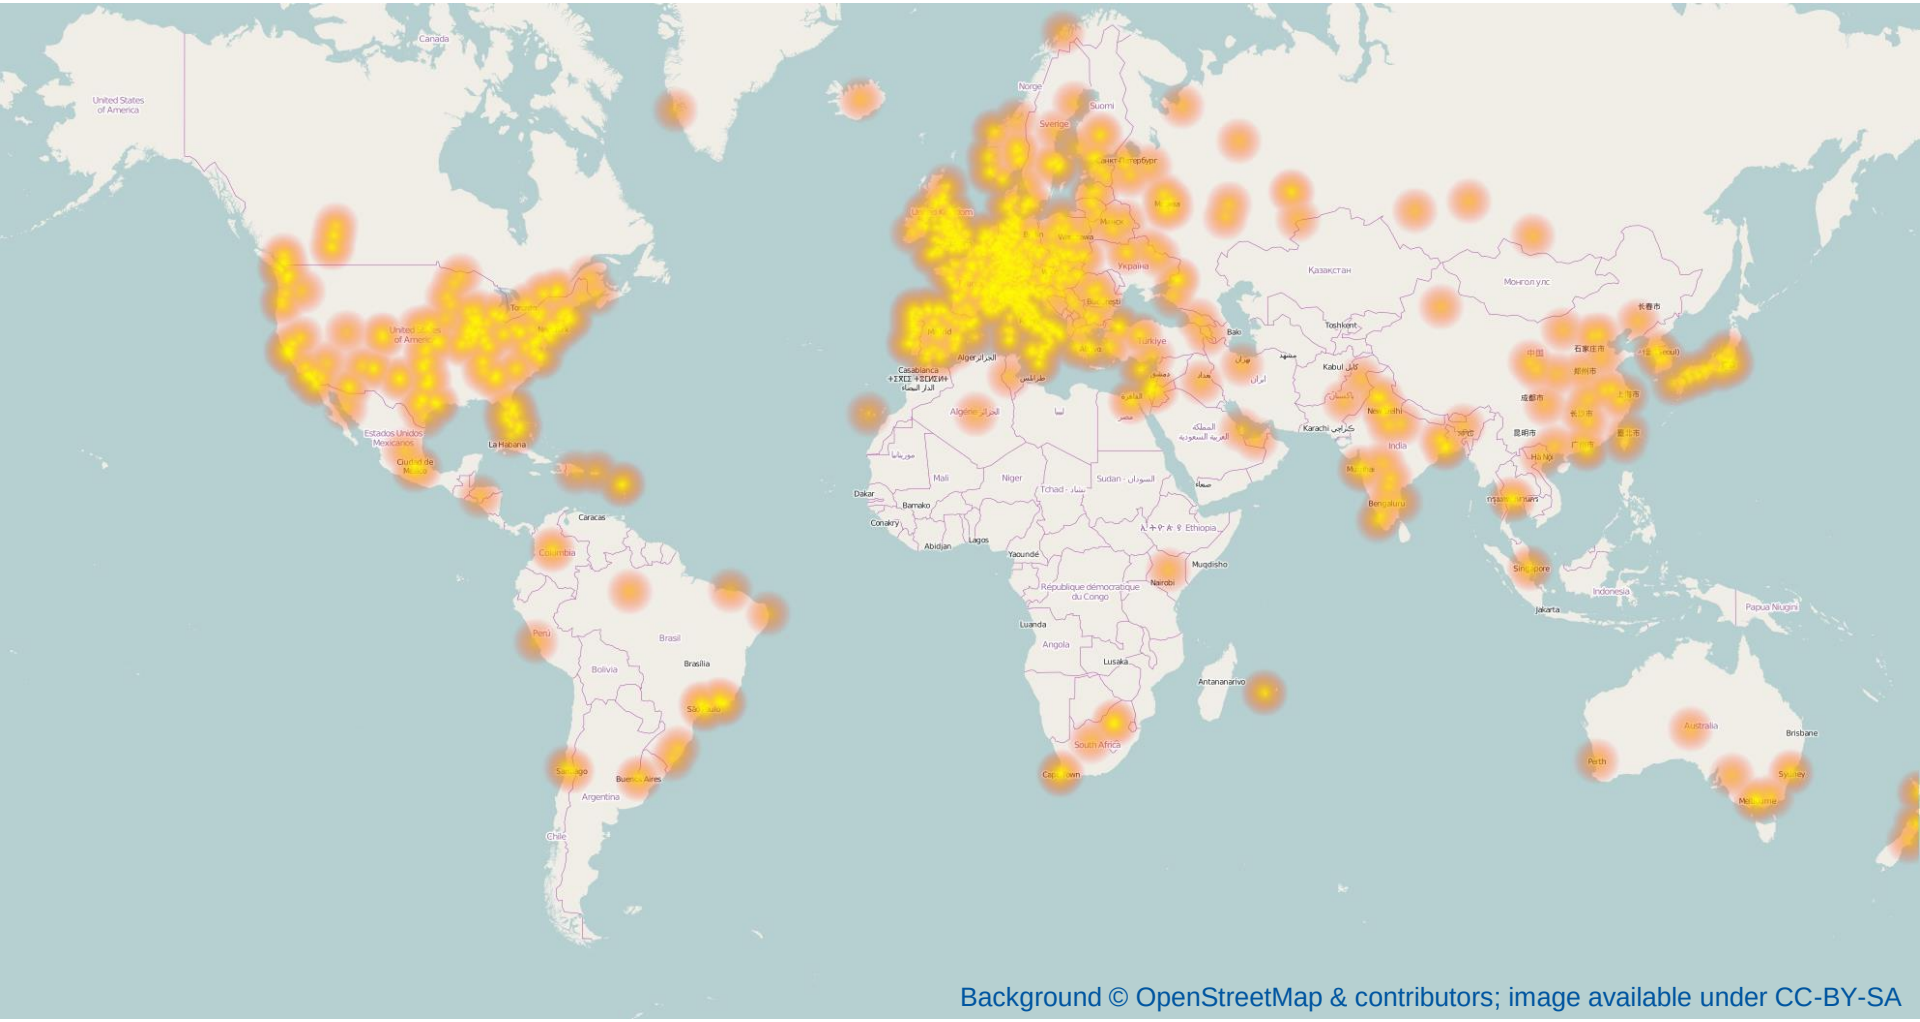

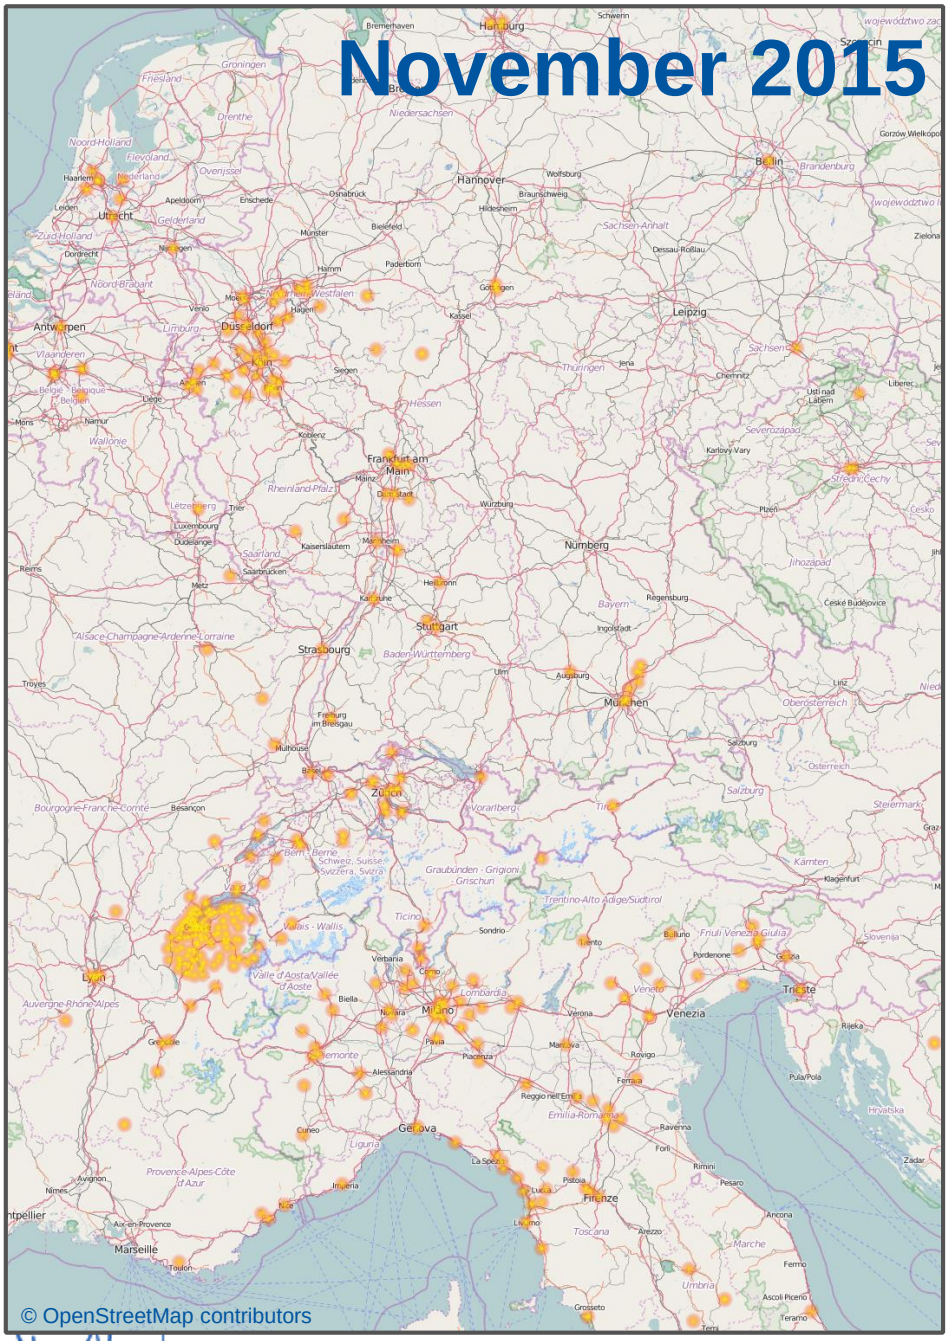
November 2015
© OpenStreetMap contributors

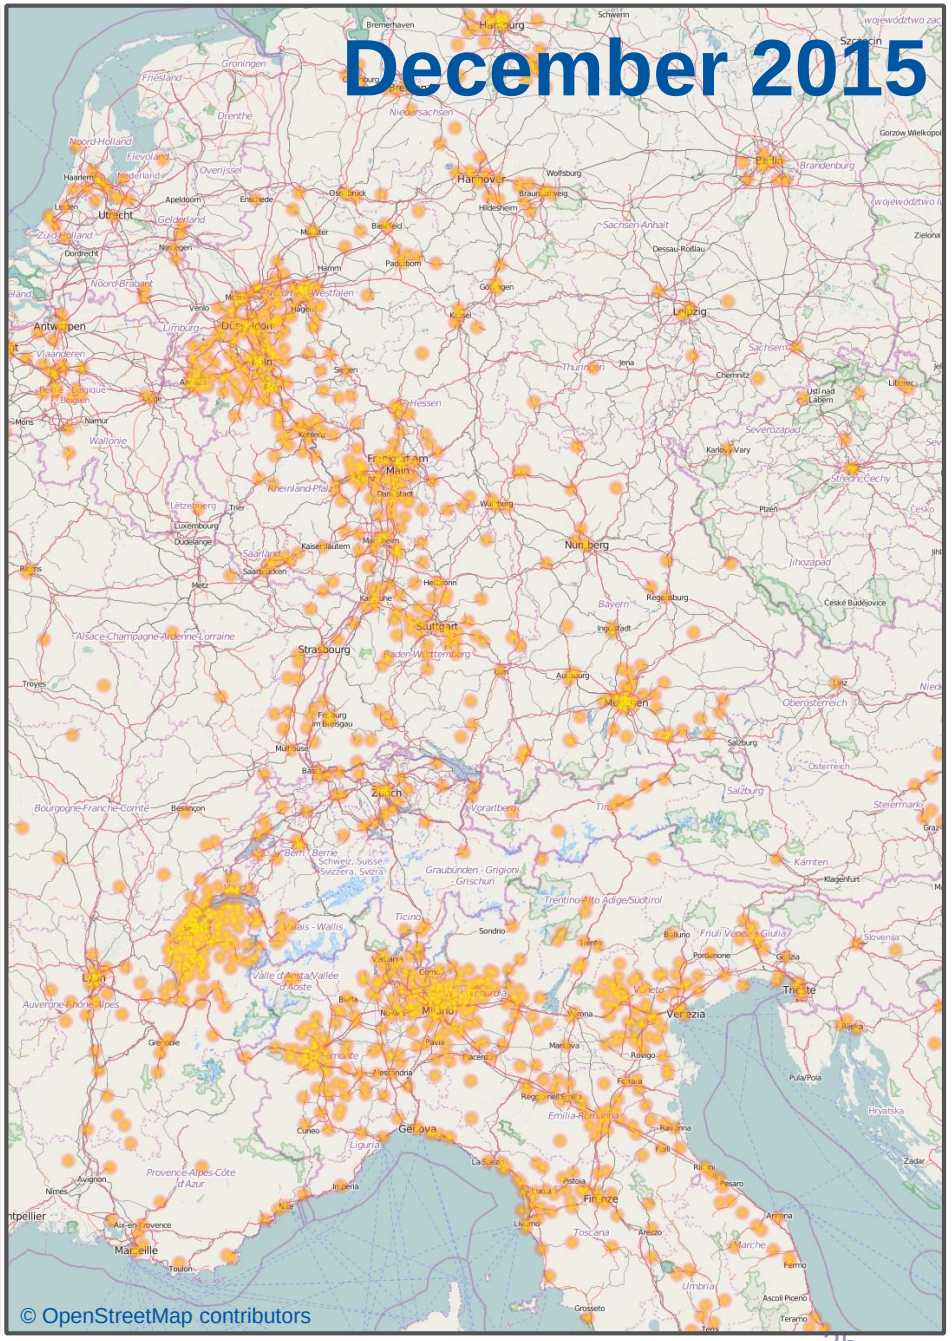
December 2015
© OpenStreetMap contributors

Usage examples

CERN Press Office

Desktop Sync Client

native client
(xrdcp)

profit from native performance in transferring 30-50GB files

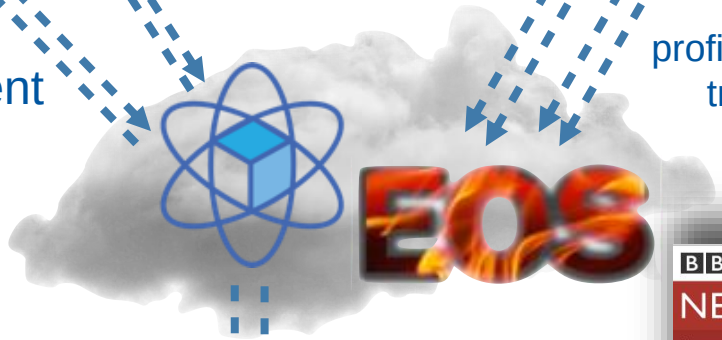

CERNBox share

CERNBox used by the **Press Office** to share immediately videos to the media for download

E-Science

United Nations Institute for Training and Research

UNOSAT
UNITAR Operational Satellite Applications Programme

RESEARCH FOUNDATION
cima
OBSERVE TO PREDICT, PREDICT TO PREVENT

CERNBox

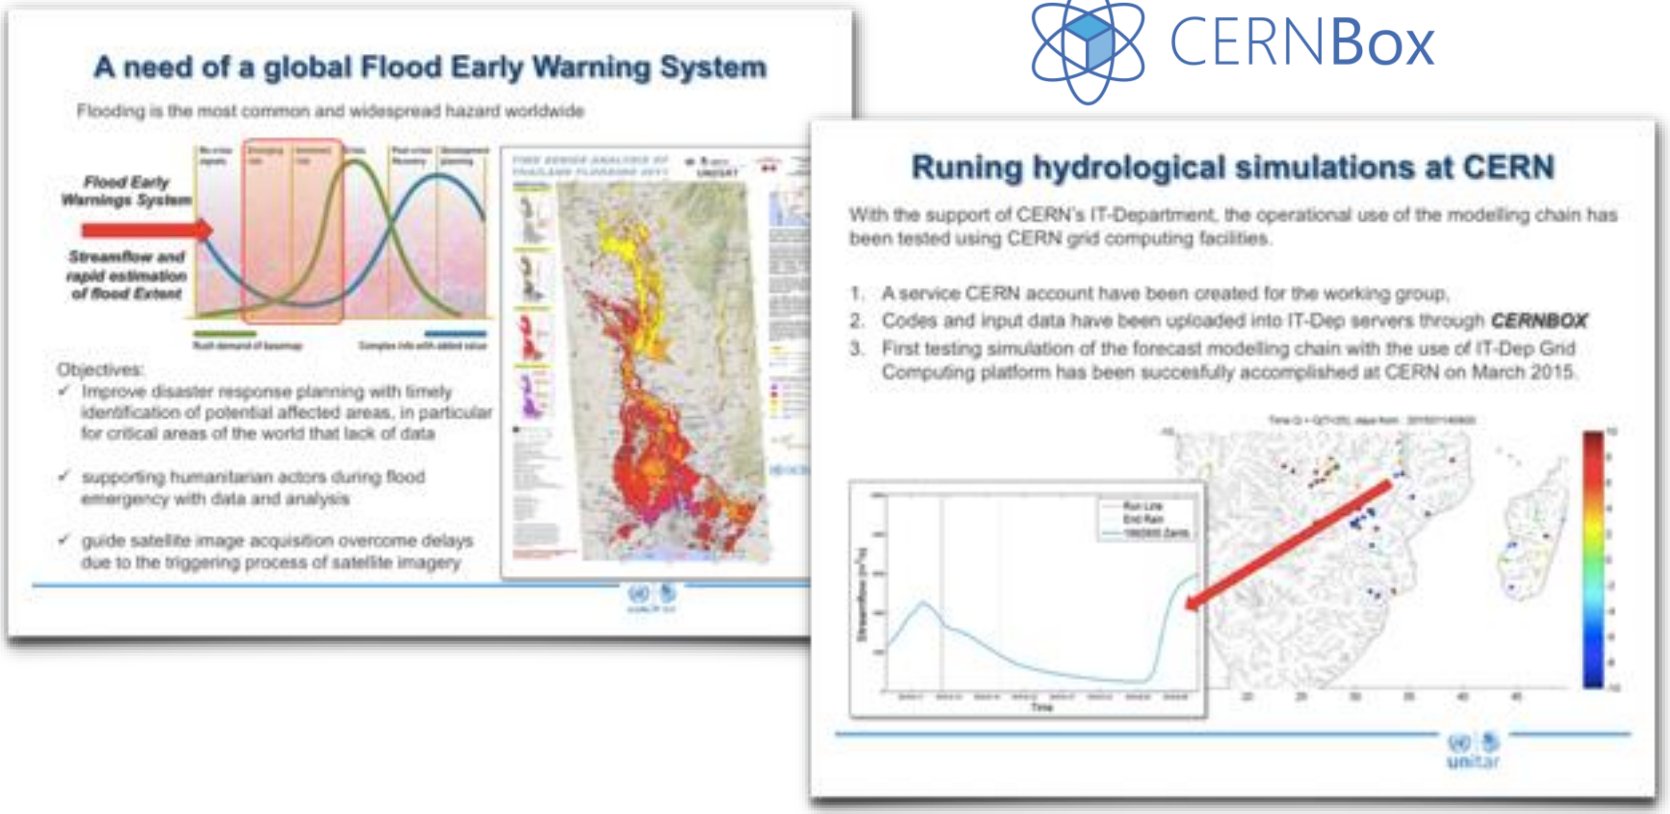

- Enable non-experts to easily use CERN Storage resources
- Powerful integration with the batch system
- Simple to share result with collaborators

Thanks to Mauro Arcorace, members of **UNITAR/UNOSAT** and **CIMA foundation** for the material provided

JRC collaboration

The Joint Research Centre (JRC)

- JRC is the science service of the European Commission
- JRC provides independent scientific support to EU policy making
- Wide usage of **Earth Observation [EO] data** as basis for research and policy support

Proposal for a "JRC Earth Observation Data Processing Platform" (JEO-DPP)

- Main focus on **satellite image** data
- Shall support existing processing workflows and environments (C/C++, Python, Matlab, Java)
- Provide different processing levels:
  - *Low-level batch processing*
  - *High-level interactive processing*
- Project timeline:
  - *Prototype development: end 2015 – mid 2017*
  - *Scaling-up in 2017/18: JRC Data Centre vs a public cloud solution*

EOS

"Earth Observation & Social Sensing Big Data Pilot Project"

- The EU ***Copernicus*** Programme with the ***Sentinel*** fleet of satellites acts as a game changer by bringing EO in the Big Data era:
  - *expected 10TB/day of **free and open** data*
  - *Requires new approaches for data management and processing*
- Pilot project launched in January 2015
- Major goal: set up a central infrastructure for storing and processing of Earth Observation and Social Sensing data at JRC

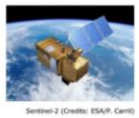

Similar to CERN SWAN

JEO-DPP processing components

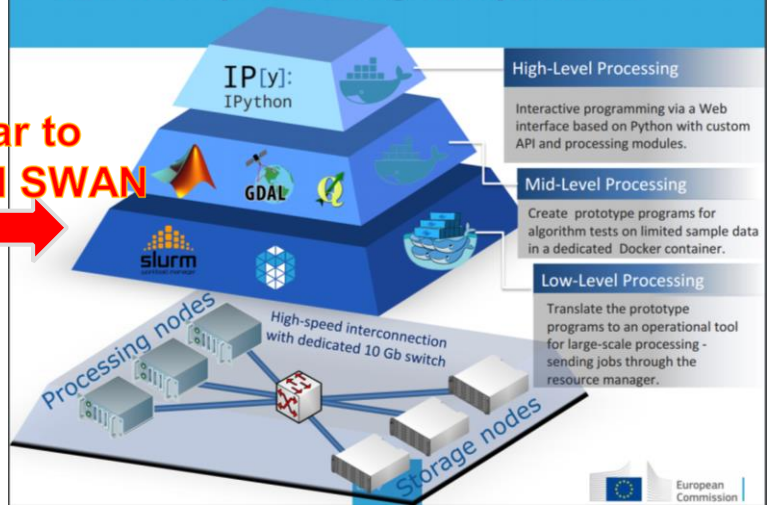

A. Burger and P.Soille (presented at the CS3 conference in Zurich – Jan 2016: cs3.ethz.ch)

R&D - EOS World-Wide Deployment

aka "exploring the 300 ms region"…

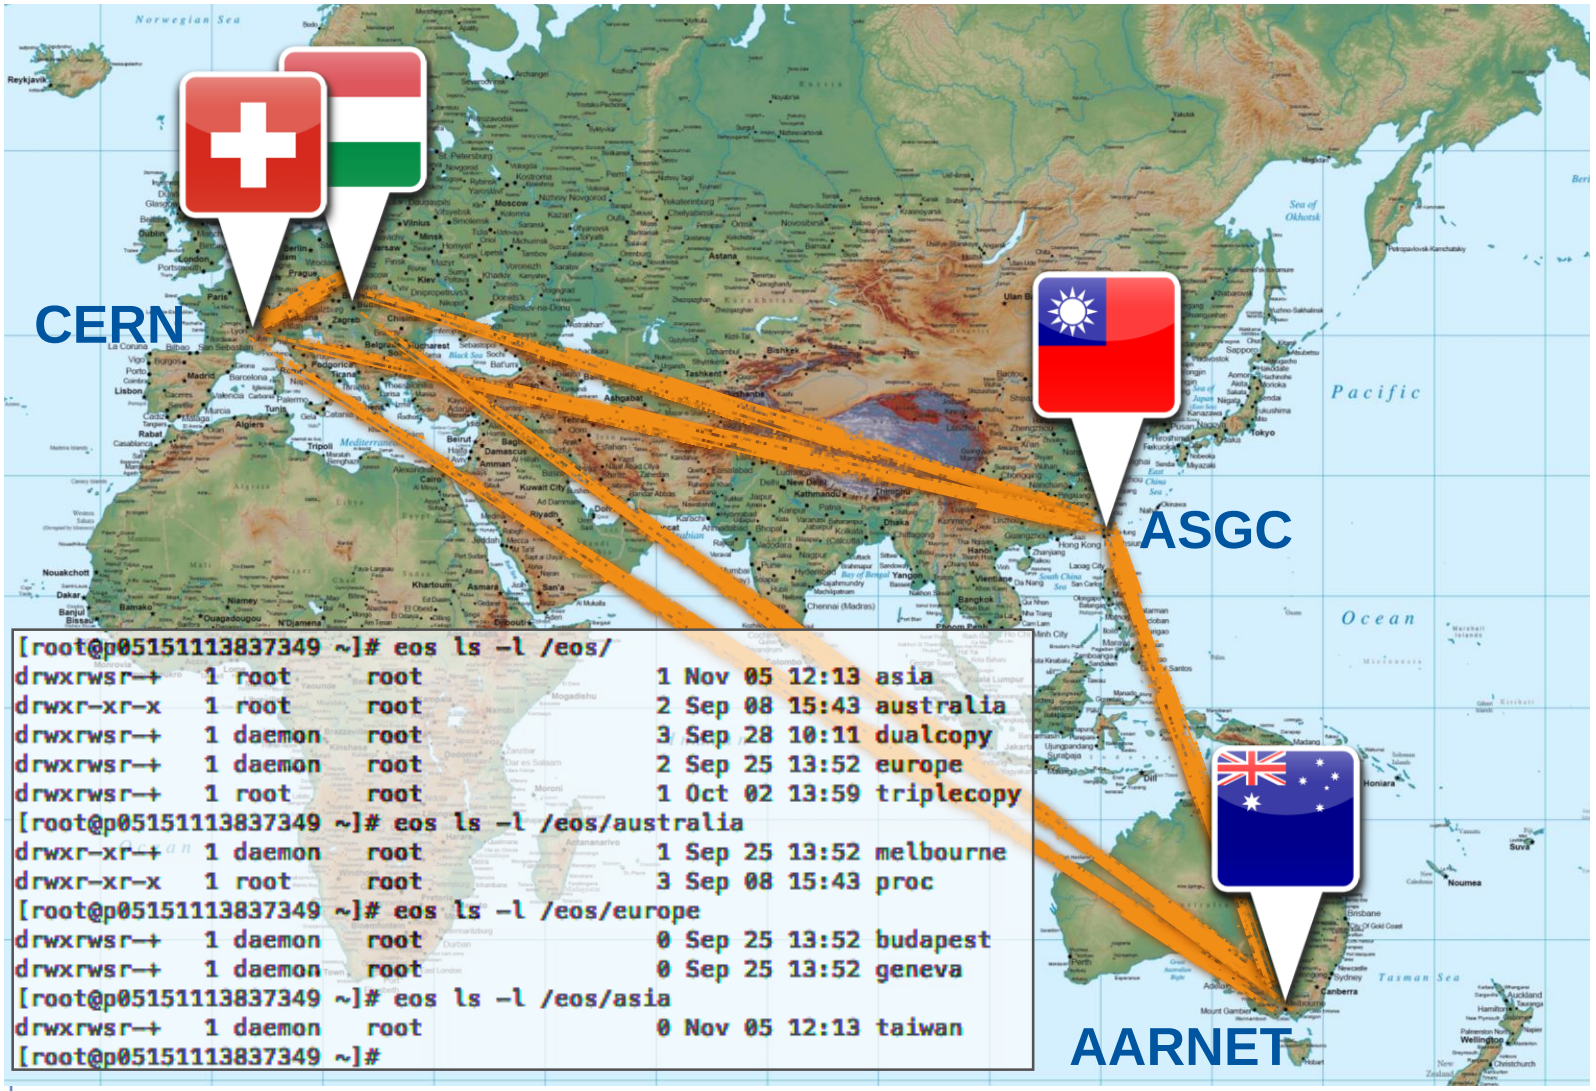

CERN data at your fingertips

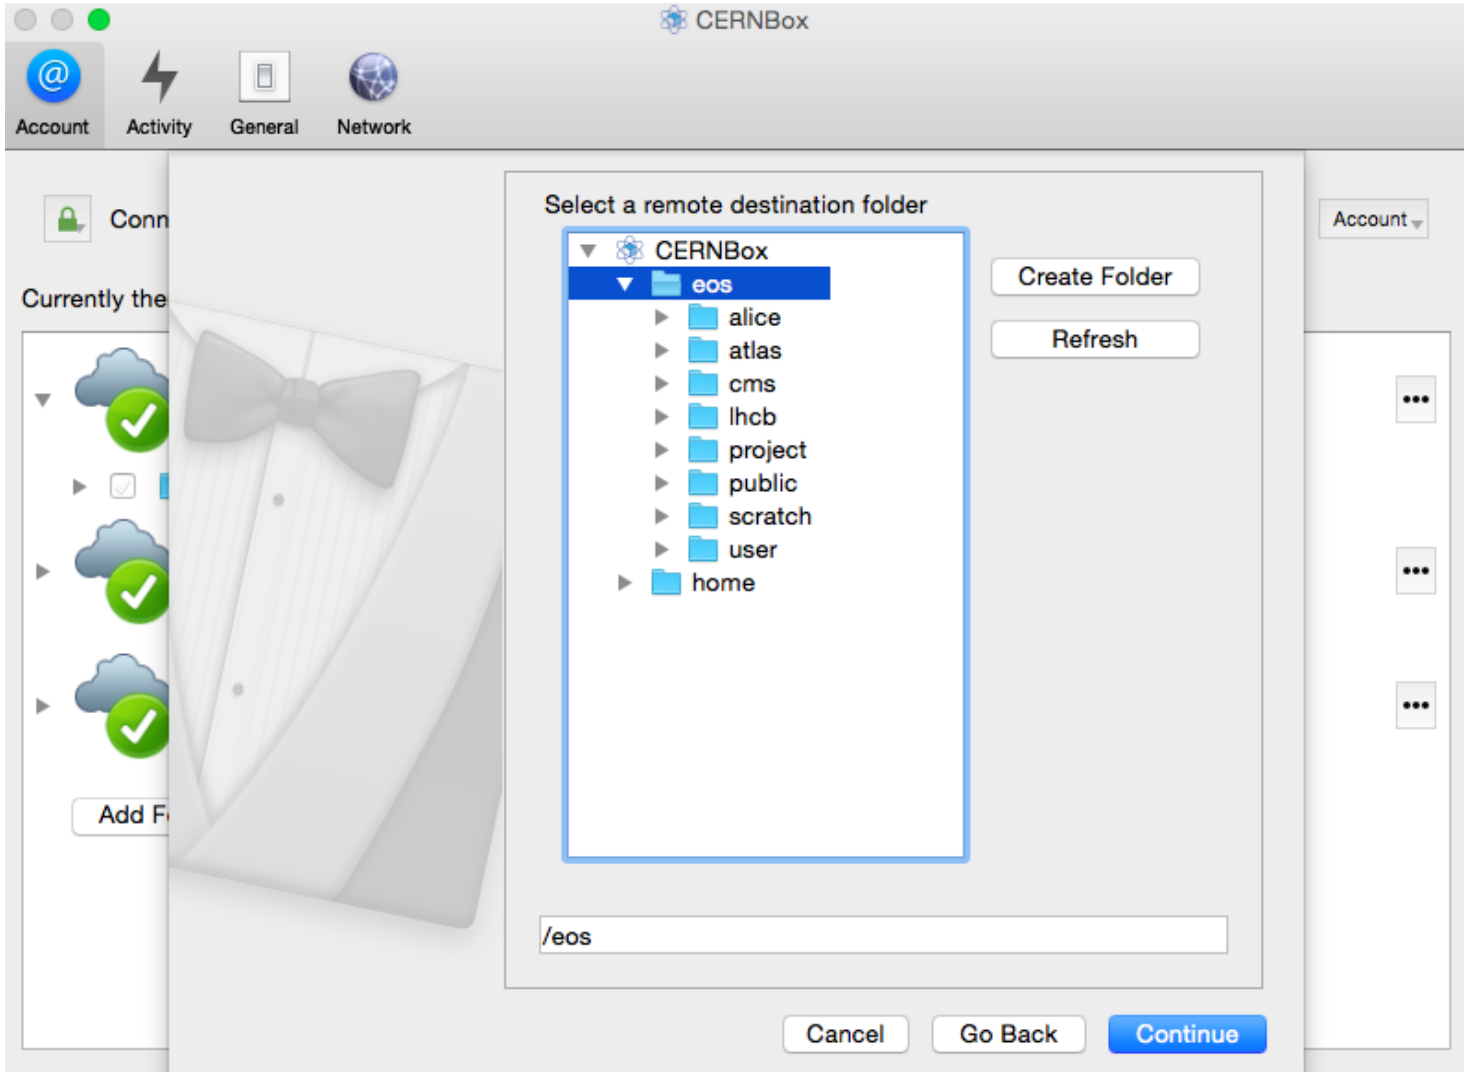

Interconnected Private Clouds for Universities and Researchers

OPENCLOUDMESH

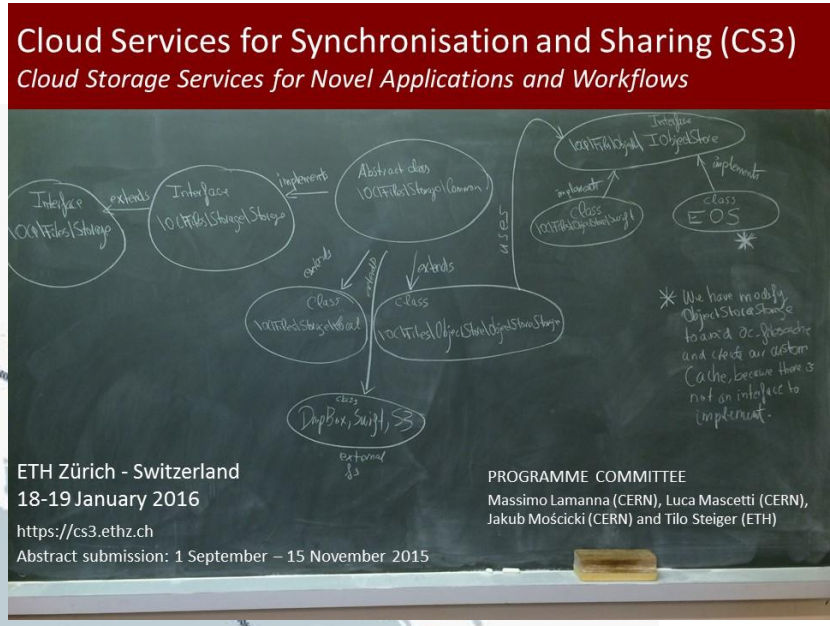

CERN
DESY Deutsches Elektronen-Synchrotron
ERCIS - European Research Center for Information Systems
ETH Zürich
Garching Computing Centre of the Max-Planck-Society
GÉANT Association
Sciebo, the 500,000 user Campuscloud
SURF, Netherlands
SWITCH, Switzerland
TU Berlin
University of Vienna

University of Florida
University of Hawaii - Manoa, College of Education

AARNet, Australia

OpenLab project

- Streamline EOS
  - More platforms
  - More “generic” installations
    - Remove CERN dependencies
  - More documentation

- Collaboration with **Comtrade**
  - 24-month project (started in July 2015)
  - http://storage.comtrade.com
  - Company based in Serbia
  - Technology interest
  - Support opportunity for EOS
    - EOS is open software

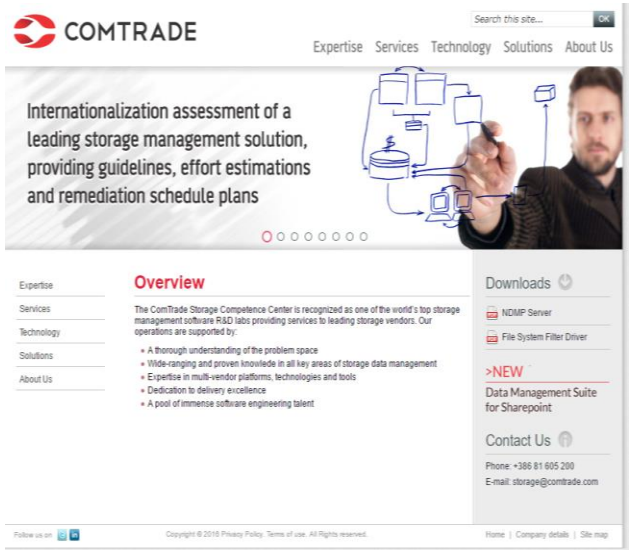

Global service for CERN users

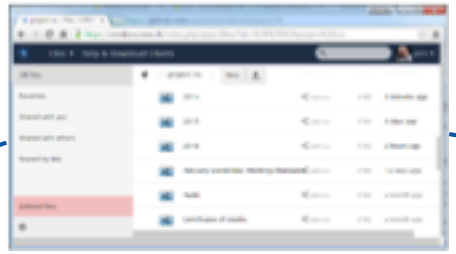

Browser and Mobile

web access

offline access

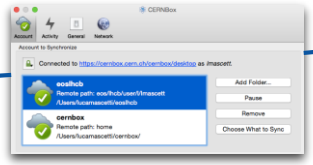

Sync Client

online access

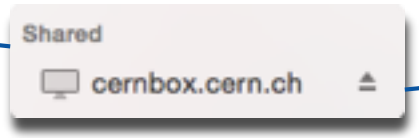

WebDAV & Fuse

Laptops, PCs
& Mobile devices

direct access

http,
webdav,
xroot

Batch Service

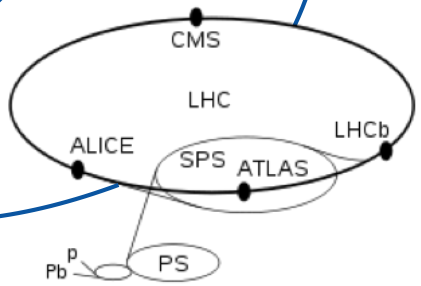

CERNBox + EOS scenarios

- end-user files and sharing
- project spaces
- physics analysis
- home directory

Notebook analysis SWAN project

with CERN Physics Department

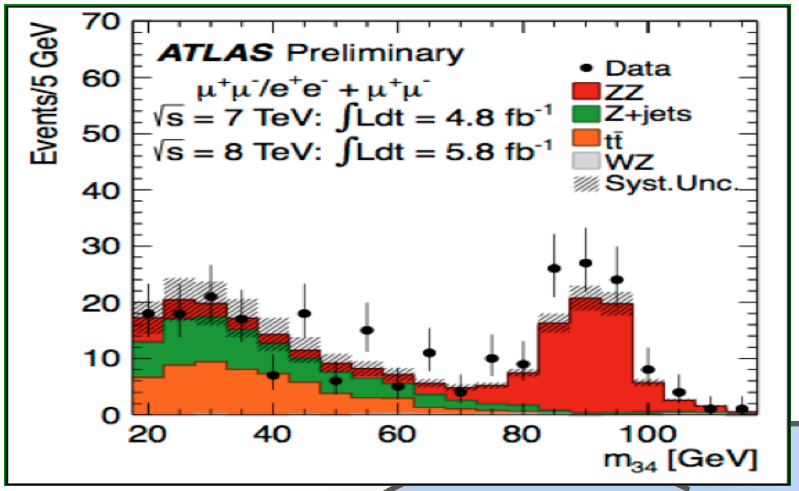

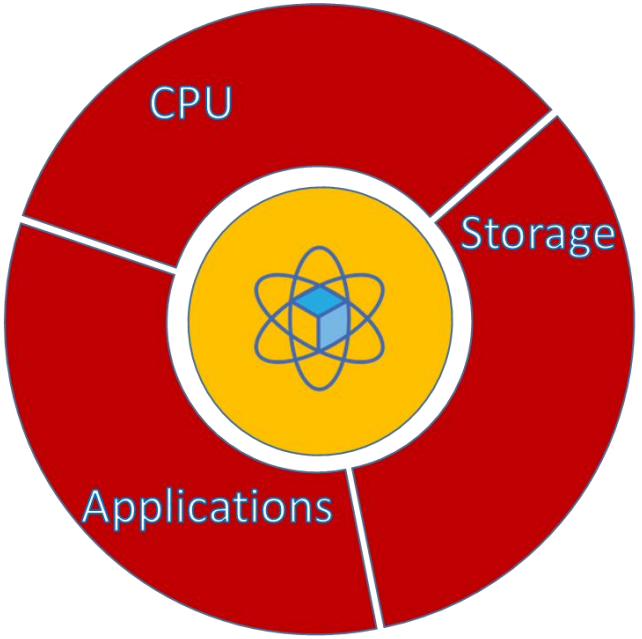

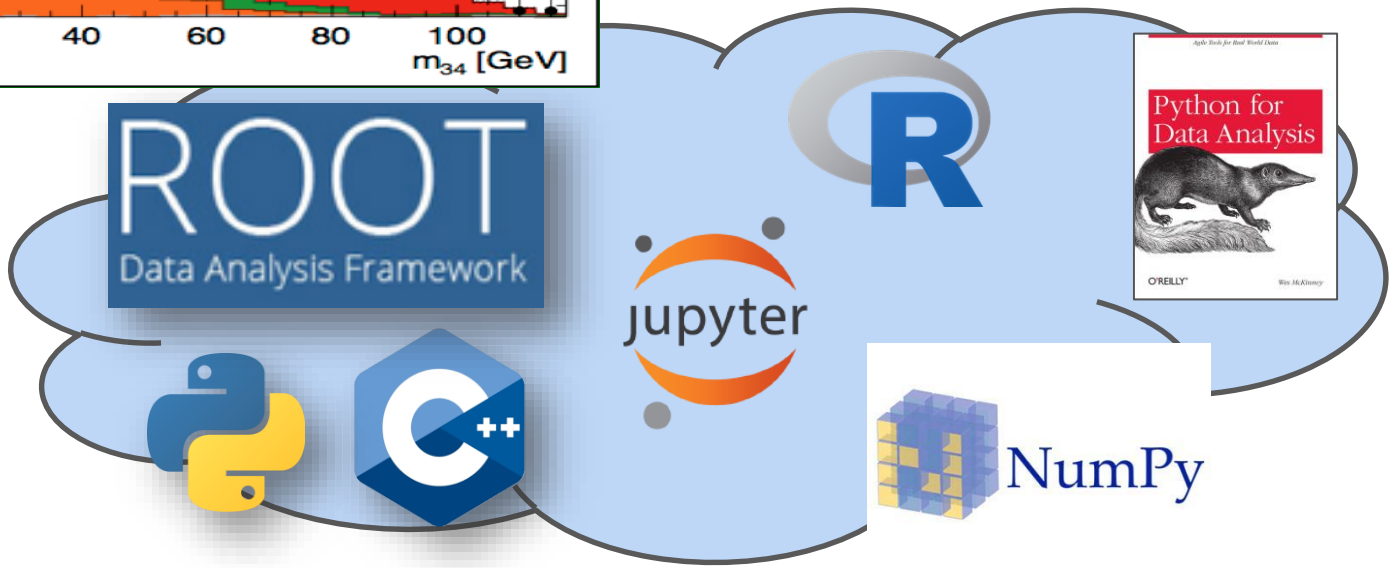

ROOT is the CERN data analysis framework: http://root.cern.ch

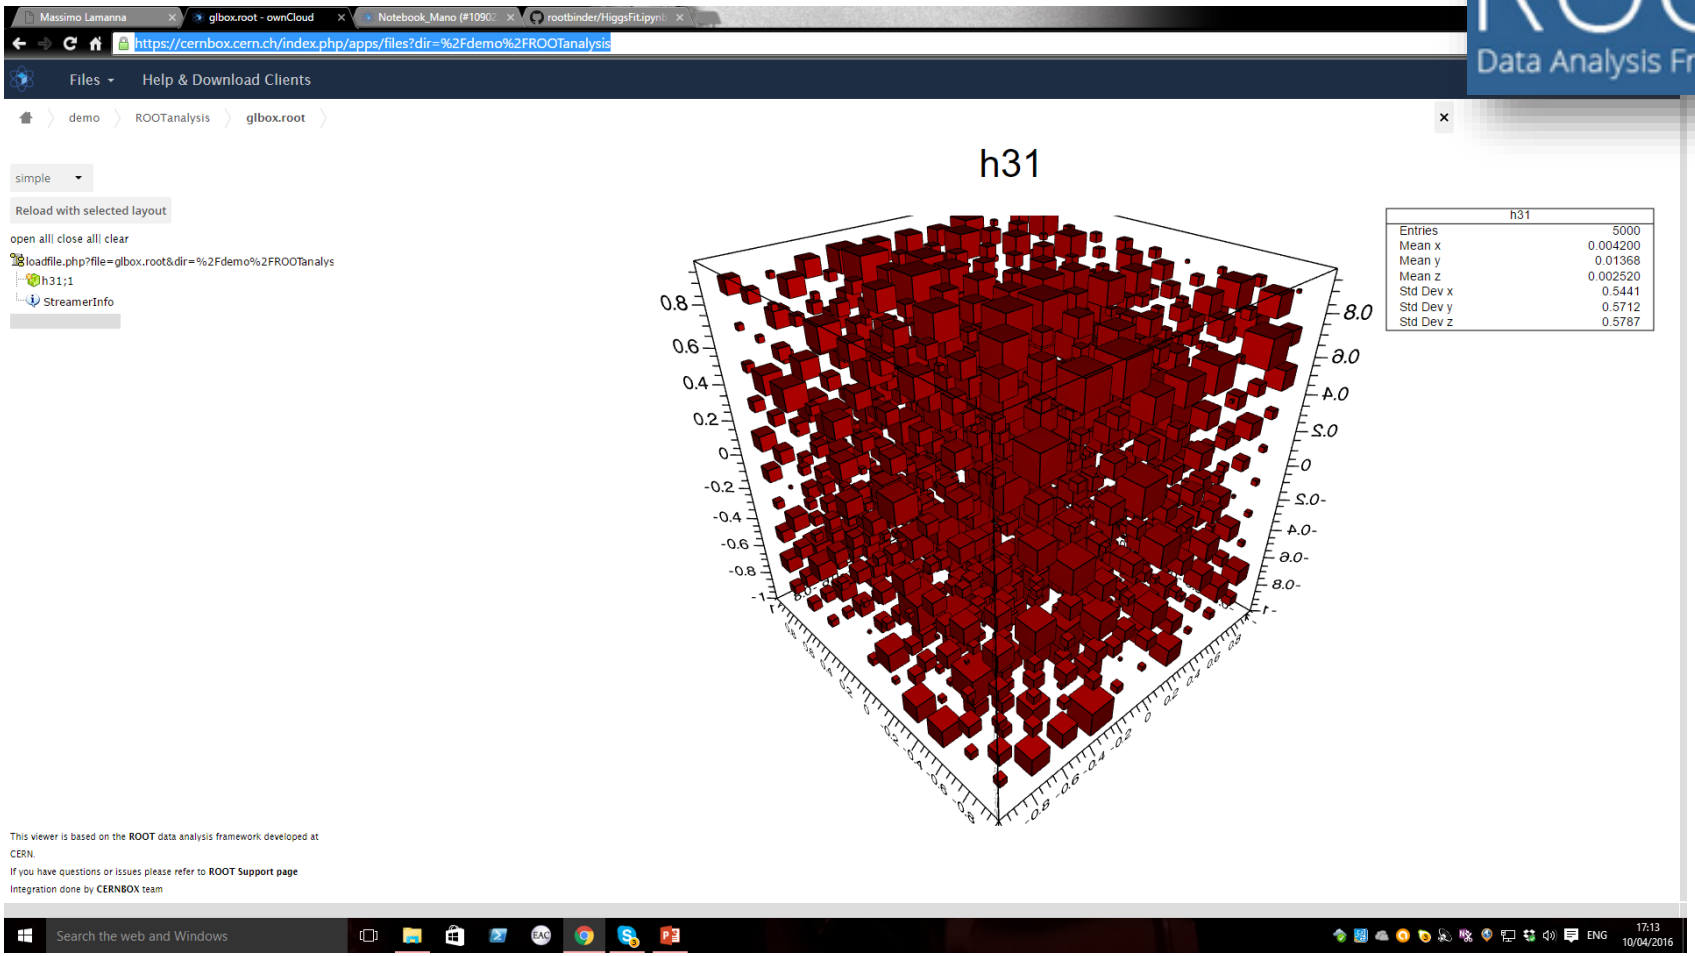

ROOT
Data Analysis Framework

Traditional data analysis → Web-based analysis

CERN

https://cernbox.cern.ch/index.php/apps/files?dir=%2Fdemo%2FROOTanalysis

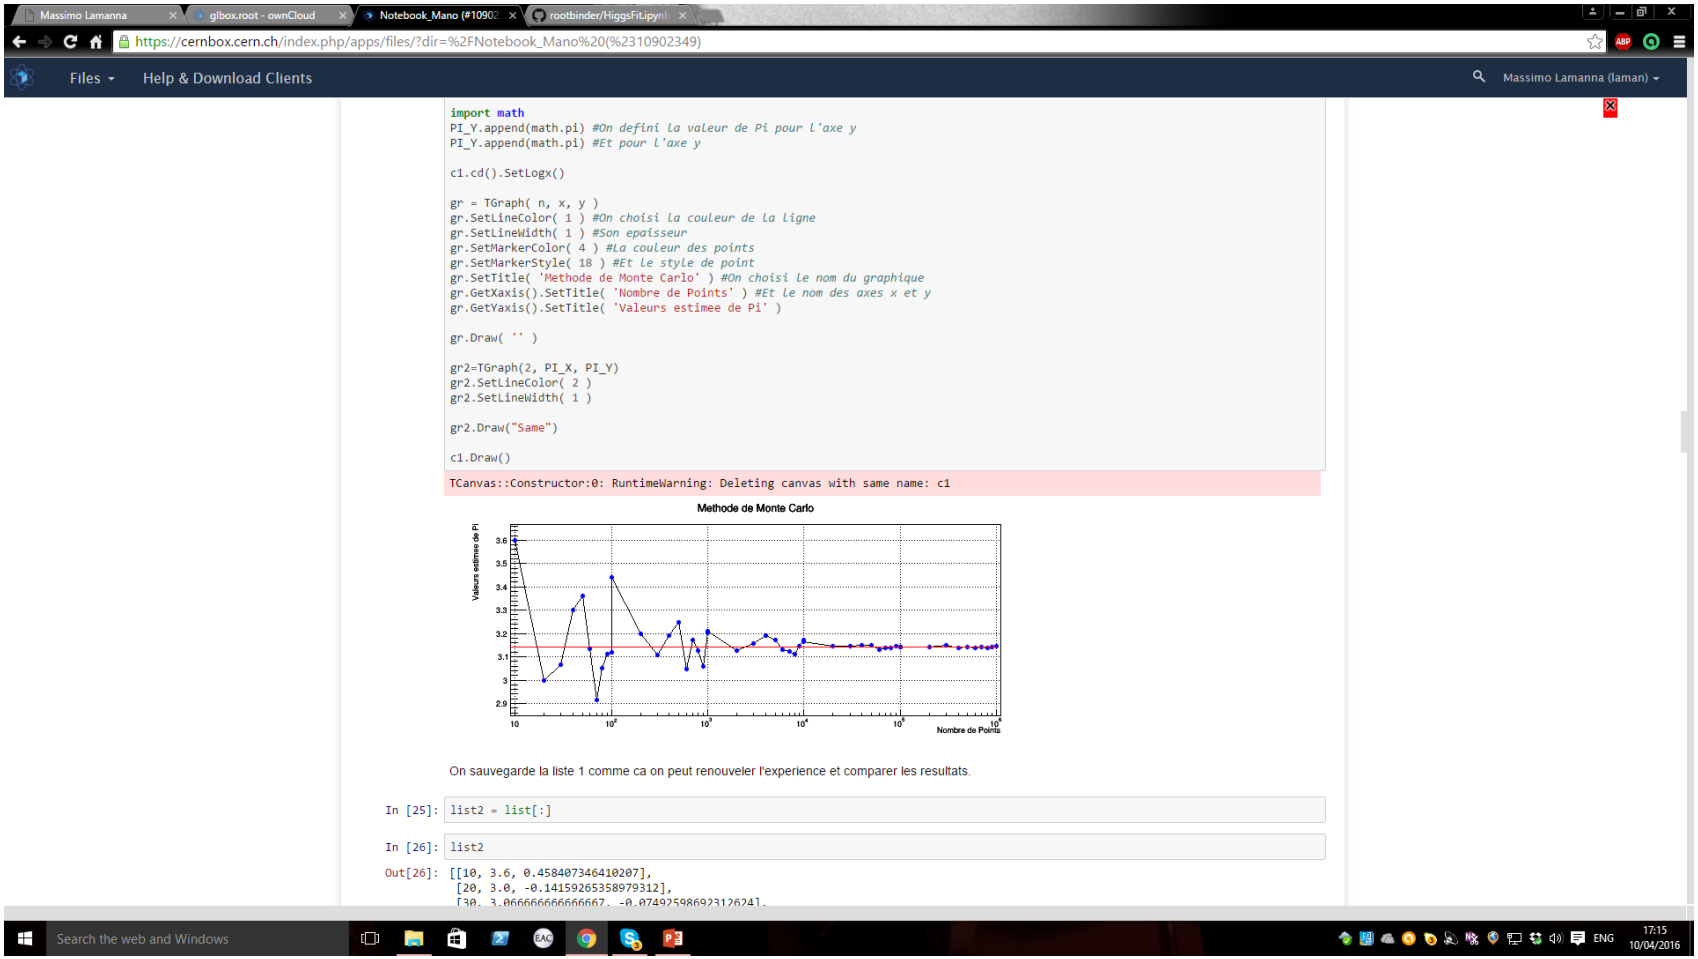

Training tool to get used with computing, mathematics, ...

CERN

Massimo Lamanna | glbox.root - ownCloud | Notebook_Mano (#10902 | rootbinder/HiggsFit.ipynb

GitHub, Inc. [US] https://github.com/cernphsft/rootbinder/blob/master/notebooks/HiggsFit.ipynb

Higgs decay to two photons

The Standard Model predicted the decay of the Higgs bosons into photons. THe process is depicted by the diagrams below:

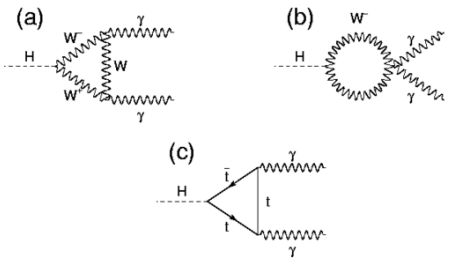

At the Large Hadron Collider, this process has been measured. This figure shows how an Higgs boson decay looks in the CMS detector:

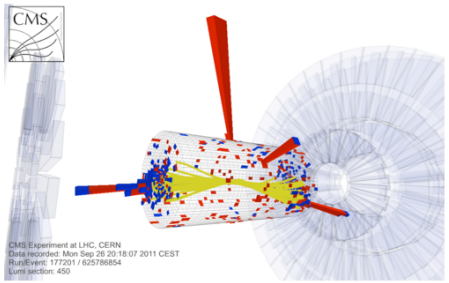

This ROOTbook illustrates a simplified fitting procedure aiming to identify the peak due to the Higgs boson decay over the exponentially falling background.

Importing input data into a ROOT file

First of all we import the input data, here simplistically stored into a text file, into a ROOT file.

In [1]:
```
TTree tree("HiggsTree","The tree cont");
auto nevt = tree.ReadFile("Hgg.txt","x");
if (nevt <= 0) {
```

Search the web and Windows

ENG 17:16 10/04/2016

Outreach to explain our research...

CERN

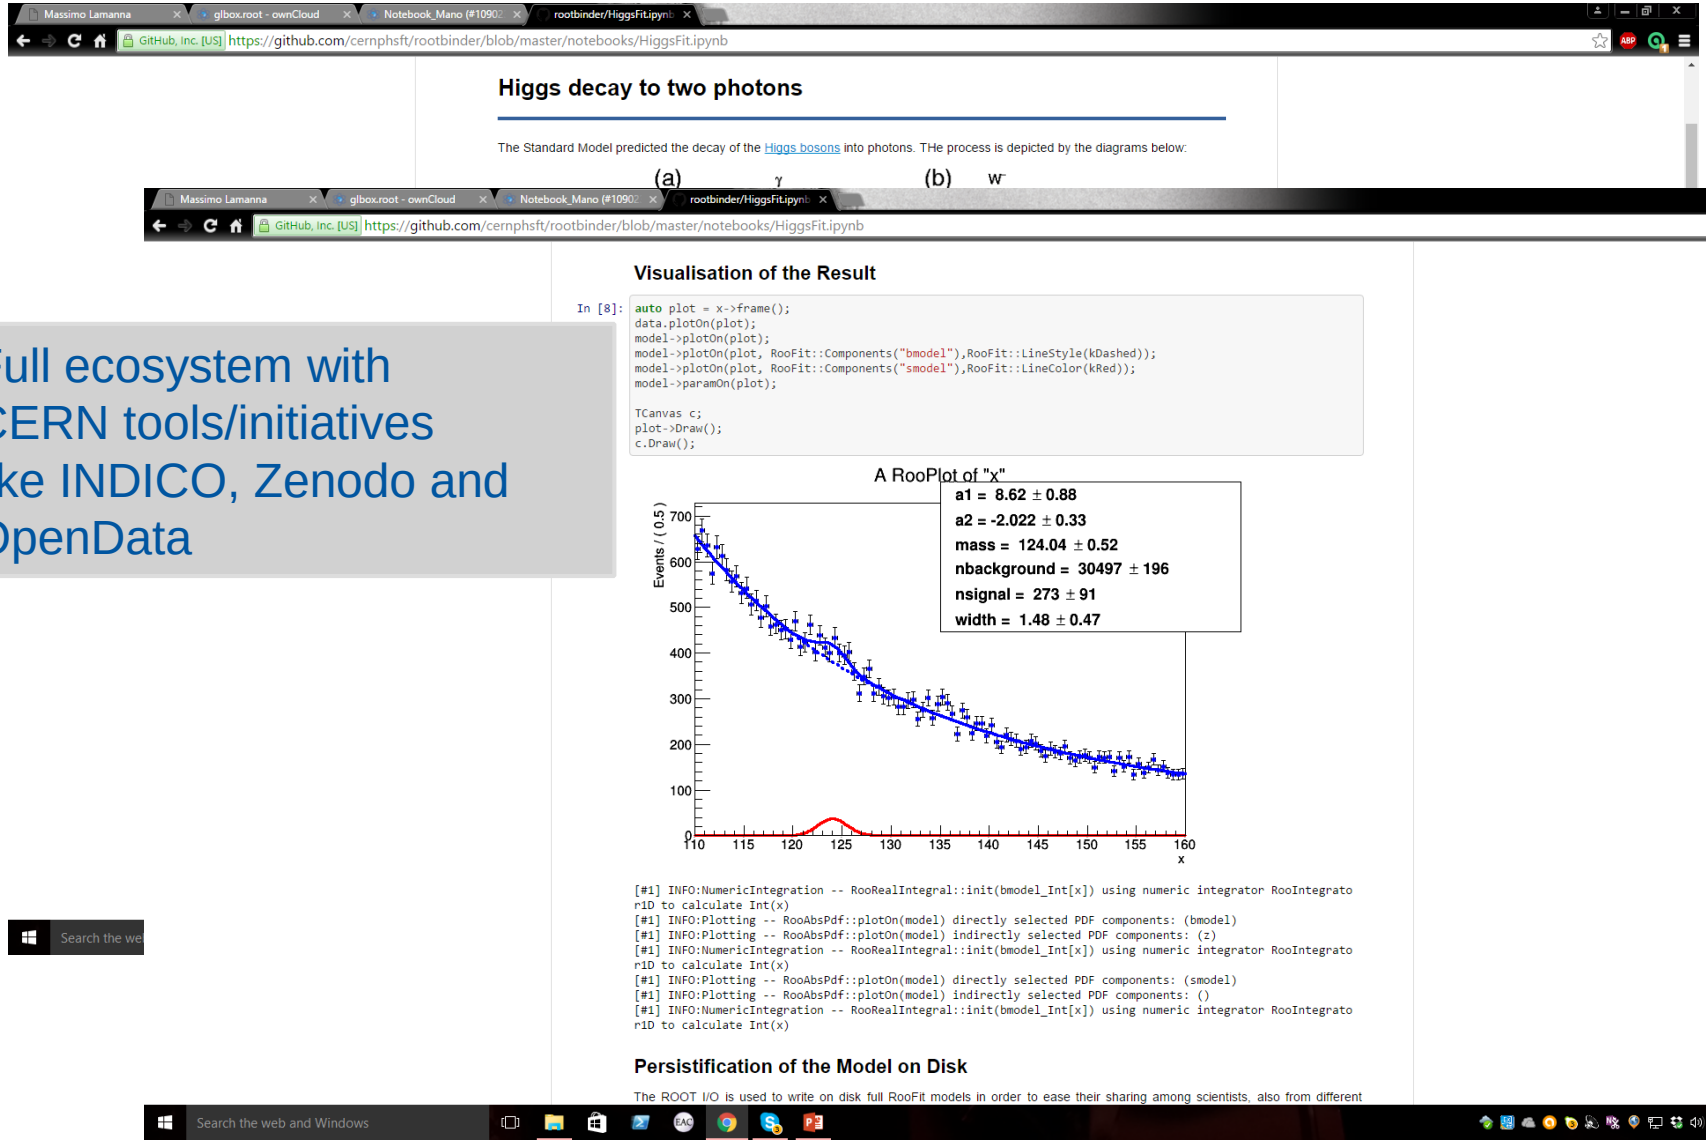

Full ecosystem with CERN tools/initiatives like INDICO, Zenodo and OpenData

Outreach to explain our research and get in touch with our data!

Summary

- Cloud storage enables new use cases
  - and new ways to work and to collaborate
- Solid foundations
  - 200 PB LHC disk infrastructure of EOS (and counting)
  - Data management experience in innovation and service provision
- CERNBox/EOS
  - Home for innovative applications
  - CS3 workshop: visit cs3.ethz.ch

DOI 10.5281/zenodo.44783

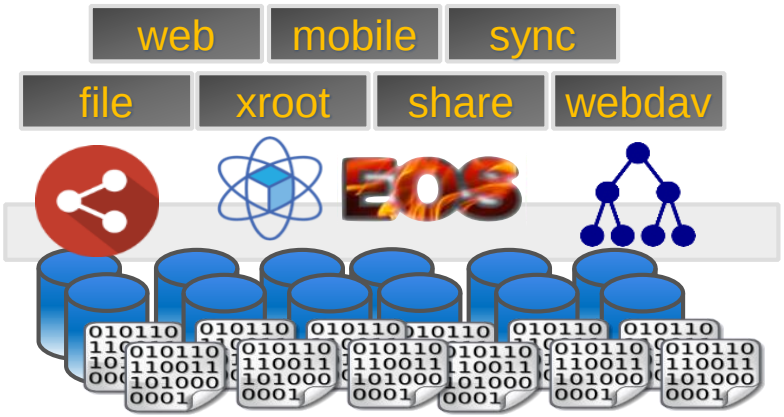

www.cern.ch